

SAN Topologies

2

This chapter describes the HP standard SAN topologies. You should review the SAN design considerations listed in the first chapter before starting the topology selection process. The design considerations enable you to generate a list of prioritized requirements for your SAN design. This list of requirements provides a basis for selecting the optimum fabric topology.

There are three approaches that you can choose when designing your SAN. You can choose to implement an HP standard SAN topology design, a subset or variation of an HP design, or you can design a custom SAN topology. Regardless of which approach you use, the final SAN design must adhere to the SAN design rules described in Chapter 3, "[SAN Fabric Design Rules](#)" and Chapter 4, "[Heterogeneous SAN Platform and Storage System Rules](#)."

Before choosing your design, you should review the HP standard SAN topology section in this chapter to get a good understanding of the important aspects of SAN implementation. HP recommends that you first consider implementing one of the HP standard topologies or a variation of one of these designs. If your requirements cannot be met by a standard design, then you can implement a customized SAN topology design—provided you follow the design rules.

Why Design Rules?

HP performs extensive qualification of all HP SAN components, including application software, operating systems, host bus adapters, Fibre Channel switches, storage systems, and storage system management appliances. A sophisticated process is used to verify the interoperability of SAN components across a wide spectrum of supported configurations, taking into consideration potential customer requirements for mixed applications, servers, operating systems, and storage systems in a single environment. In order to ensure that a new SAN installation will function properly, certain guidelines must be followed. These guidelines are the result of actual laboratory testing—not theoretical projections of “what should work”—and reflect the designed-in capabilities of the various software and hardware components that are used in SAN storage systems.

The guidelines are captured in rule form to make it easier to design a SAN that will work properly and will be supported by HP. Chapter 3, "[SAN Fabric Design Rules](#)" and Chapter 4, "[Heterogeneous SAN Platform and Storage System Rules](#)", define the configuration rules. For additional information on operating system HBA/driver/firmware/software support, contact your HP field representative.

Switch Product Lines

HP supports three product lines of Fibre Channel switch products that may be used to build SAN fabrics. Each product line provides certain advantages that apply to specific applications. For more information on specific switch models, refer to the section [Switch Model Selection Guidelines](#) and the Fibre Channel switch product information on the Storage SAN Infrastructure web page accessible through www.hp.com.

The B-Series product line includes a wide range of Fibre Channel switches, described as "SAN switches" and "Core and Director switches." A partial list of products in this family includes the HP StorageWorks SAN Switch 4/32 and the HP StorageWorks SAN Director 2/128. This product line includes switches with 8, 16, 32, 64, and 128 ports, including both full-function and entry-level models. The HP StorageWorks Core Switch 2/64 and SAN Director 2/128 have a high level of internal redundancy. The B-Series product line also includes the HP Multi-Protocol Router (MP Router). The MP Router supports Fibre Channel routing and FCIP.

The C-Series product line includes the Cisco MDS 9506 and 9509 Multilayer Directors and the MDS 9216/9216A/9216i, 9120, and 9140 Multilayer Fabric Switches. The MDS 9509 is supported with 224 ports and MDS 9506 is supported with up to 128 ports per chassis. MDS 9509 and 9506 offer 7 and 4 modular slots respectively that can be populated with 16-port, 32-port, or Cisco MDS 9000 4/8-Port IP Storage Services module or Cisco MDS 8000 14/2-Port Multiprotocol Services module. The MDS 9216/9216A/9216i have 2 slots -- one slot is a fixed configuration with either a 16 fibre channel port module in the 9216/9216A or a 14 fibre channel ports and 2 ports of GE (Gigabit Ethernet) in the 9216i switch. The second - an expansion slot that supports either a 16 or a 32 port card, for 32 or 48 ports in total. The second slot can also be populated with a Cisco MDS 9000 4/8-Port IP Storage Services module supporting FCIP and iSCSI via 4/8 ports of GE (Gigabit Ethernet).

The MDS 9100 Series of multilayer fabric switches is a fixed-configuration 1U platform consisting of the MDS 9120 and MDS 9140 switches supporting a fixed 20 port and 40 ports respectively. The MDS 9100 Series supports the same multilayer intelligent networking services as the MDS 9500 Series and MDS 9200 Series including the same command line interface (CLI) and embedded Fabric Manager suite.

The M-Series product line includes a wide range of Fibre Channel switches described as "Directors" and "Edge switches." A partial list of products in this family includes the HP StorageWorks Director 2/140 and the HP StorageWorks Edge Switch 2/32. This product line includes switches with 12, 24, 32, 64, and 140 ports, and all models use the same version of internal microcode. The HP StorageWorks Director 2/64 and 2/140 have a high level of internal redundancy.

The switch model numbering convention is the same in B-Series and M-Series product families. The number preceding the slash indicates the highest speed at which the switch ports can operate, measured in Gbps, and the number following the slash indicates the number of ports on the switch. The HP StorageWorks SAN Switch 2/16 is a 2 Gbps switch with 16 ports, and the HP StorageWorks Edge Switch 2/32 is a 2 Gbps switch with 32 ports.

For B-Series switches, entry level models are indicated by the suffix "-EL", "V", or "N" in the product name.

All HP 2 Gbps Fibre Channel products implement ports that auto negotiate their signaling speeds. Each pair of ports uses the lower of the supported speeds, so if a 2 Gbps port is connected to a 1 Gbps port, they both run at 1 Gbps in each direction. This applies to pairs of ports that are directly connected together. The speed of a remote port does not affect the local port speed, because speed matching is done within the switches in the fabric.

Switch Product Line Interoperability

For new SAN deployments HP recommends you utilize switch models from a single product line exclusively. To meet the needs of customers desiring a mix of switch models from the different product lines however, HP does support three levels of SAN fabric interoperability. Specifically, HP supports:

- Within a multi-fabric SAN, one fabric with all B-Series switches and another fabric with all M-Series switches. This is referred to as a "dual heterogeneous SAN fabric"
- Within a single fabric, M-Series Director and Edge switch models intermixed with B-Series SAN switch models. This is referred to as an "interoperable heterogeneous SAN fabric"
- Within a single fabric, C-Series Director and Fabric switch models intermixed with B-Series switch models.

Refer to [Heterogeneous/Interoperable SAN Fabrics](#), page 98 for the specific fabric interoperability rules and supported switch models for these levels of interoperability.

Switch Model Selection Guidelines

Each of the three HP Fibre Channel switch product lines offers a range of switch models including core or director, and SAN, edge or fabric switch types. For the purposes of selection guidelines in this section of the guide, the terms “core” and “edge” are used generically in reference to switch placement in a core to edge fabric design. Elsewhere in this guide and in other HP literature, the use of these terms more formally within a product name describes very specific switch models such as “Core Switch 2/64”, “Director 2/140” or “MDS Fibre Channel Director”, and “SAN Switch 2/16”, “Edge Switch 2/16” or “MDS Fibre Channel Fabric Switch”.

For each switch product line, the general recommendations and guidelines presented are based on the combination of a number of factors such as switch cost, scalability, and availability features. Specific customer requirements may rank some factors higher than others in importance and need to be considered when selecting the appropriate switch model for a specific SAN implementation. Refer to the switch model features table for each product line for specific recommendations based on certain switch model features.

Note: The switch model selection information presented in this section provides general guidelines for usage, and is not meant be a substitute for a thorough architectural level topology design process. Refer to other sections of this guide for information about determining the optimal fabric topology design prior to selecting specific switch models.

B-Series Switch Model Selection Guidelines

The following two tables show the recommended usage of each B-Series switch model for core or edge switch placement in a fabric design. Refer to Table 4 for specific switch model features.

Table 2: B-Series switch model usage as a core switch

B-Series Core Switch Model Selection					
	1-96 User Ports	97-224 User Ports	225-500 User Ports	501-728 User Ports	728-1280 User Ports
Core Switch 2/64, SAN Director 2/128	Excellent	Excellent	Excellent	Excellent	Excellent
SAN Switch 2/32, SAN Switch 4/32	Excellent	Very good	Good	Not recommended as a core switch	Not recommended as a core switch
SAN Switch 2/8, 2/8-EL, 2/16, 2/8V, 2/16V, 2/16N, and 2/16N FF	Good	Good	Good	Not recommended as a core switch	Not recommended as a core switch
SAN Switch 8, 16	Good	Good	Not Recommended	Not recommended as a core switch	Not supported
Note: The SAN Switch 8 and SAN Switch 16 are only recommended as core switches when connected to other 1 Gbps SAN switches					

Table 3: B-Series switch model usage as a SAN switch

B-Series SAN Switch Model Selection					
	1-96 User Ports	97-224 User Ports	225-500 User Ports	501-728 User Ports	728-1280 User Ports
Core Switch 2/64, SAN Director 2/128	Excellent	Excellent	Excellent	Excellent	Excellent
SAN Switch 2/32, 4/32	Excellent	Excellent	Excellent	Excellent	Excellent
SAN Switch 2/8, 2/8-EL, 2/16, 2/8V, 2/16V, 2/16N, and 2/16N FF	Excellent	Very Good	Very Good	Good	Good
SAN Switch 8, 16	Very Good	Very Good	Good	Good	Not Supported

Table 4: B-Series switch model features

	Redundant Power, Hot Swap Power	Redundant Cooling, Hot Swap Cooling	Redundant Control Processor	Non-disruptive Code Activation	Non-disruptive Port Expansion	High Availability (Redundant Active Components)
SAN Director 2/128	Yes, Yes	Yes, Yes	Yes	Yes	Yes	Yes
Core Switch 2/64	Yes, Yes	Yes, Yes	Yes	Yes	Yes	Yes
SAN Switch 2/32, 4/32	Yes, Yes	Yes, Yes	No	Yes	No	No
SAN Switch 2/16	Yes, Yes	Yes, Yes	No	No	No	No
SAN Switch 8, 16, 2/8, 2/8-EL	No, No	Yes, No	No	No	No	No

B-Series Multi-Protocol Support

B-Series switches are supported for Fibre Channel Routing and FCIP with the B-Series MP Router, and for iSCSI and FCIP through the SR2122-2 storage router. Refer to Chapter 3, "[SAN Fabric Design Rules](#)" and Chapter 7, "[SAN Extension](#)" for more information.

B-Series Switch Model Features

- **Investment protection** Upgrading from EL switches to full fabric functionality without disruption in a production environment.
- **Compatibility** Plug and play installation with currently installed B-Series switch ports.

power pack Option License bundled software features include:

- **Advance Performance Monitor** - Provides complete resource utilization analysis on a fabric-wide basis.
- **Advance WebTools** - A browser-based application for managing B-Series switches, WebTools simplifies management by enabling administrators to configure, monitor, and manage switch and fabric parameters from a single online access point.

- **Advanced Zoning** - Providing secure access control over fabric resources. Both port and WWN zoning is hardware enforced at the frame level.
- **Extended Fabrics** - Enable Fibre Channel SAN connectivity at distances up to 100 km. Extended Fabrics applications improve disaster recovery operations and help ensure business continuance.
- **Fabric Watch** - Fabric Watch proactively monitors the health and performance of switches and the SAN fabric from a central point.
- **ISL Trunking** Inter-Switch Link Trunking is an ASIC based feature which combines multiple links between switches forming a single logical ISL with a total bandwidth of 8 Gbps. ISL Trunking enhances performance via dynamic load balancing frames across links.
- **Remote Switch** Remote Switch enables B-Series switches to create one logical SAN that spans remote fabrics at unlimited distances. SAN to WAN connectivity enables one logical view of a SAN in which all components appear as local devices.

Additional software features and applications:

- **Secure Fabric OS** Secure Fabric OS is a comprehensive security solution for B-Series SAN fabrics. Secure Fabric OS provides flexible security and policy based administration that protects data from unauthorized access and corruption. Licensed separately.
- **Fabric Manager** Fabric Manager is a host based application which provides administrators with the ability to simplify B-Series SAN management. The Fabric Manager application is sold separately.

C-Series Switch Model Usage

The following table shows the recommended usage of each C-Series switch model for core or edge switch placement in a fabric. Refer to Table 6 for specific switch model features.

Table 5: C-Series switch model usage as a function of the fabric size

C-Series Director Switch Model Selection			
	Up to 48 Total Ports	48 to 224 Total Ports	224-3000 Total Ports
Recommended Topology	Single Switch	Single Director	Core/Edge
MDS 9509 Director	Good ^{1,2}	Excellent	Excellent / Core Director
- 16-port Module	Storage and ISL	Storage and ISL	Storage and ISL
- 32-port Module	Host & Tape	Host & Tape	Host & Tape
MDS 9506 Director	Very Good ²	Excellent / Up to 128 ports	Excellent / Core Director
- 16-port Module	Storage and ISL	Storage and ISL	Storage and ISL
- 32-port Module	Host and Tape	Host and Tape	Host and Tape
MDS 9216/9216A/9216i Fabric Switch	Excellent Up to 48 ports (46 ports for 9216i)	Excellent Edge Switch	Very Good / Edge Switch

Table 5: C-Series switch model usage as a function of the fabric size

C-Series Director Switch Model Selection			
- 16-port Module	Storage and ISL	Storage and ISL	Storage and ISL
- 32-port Module	Host and Tape	Host and Tape	Host and Tape
MDS 9120 Fabric Switch	Excellent / Up to 20 ports	Excellent Edge Switch	Excellent / Edge Switch
MDS 9140 Fabric Switch	Excellent / Up to 40 ports	Excellent Edge Switch	Excellent / Edge Switch

1. Excellent if planning to scale to larger port count in the near future
2. Excellent if ultra high availability at the device level is a requirement

C-Series Switch Model Features

All C-Series products come with the following standard features:

- Non-blocking architecture using Virtual Output Queuing
- VSAN (Virtual Storage Area Network) for deployment of secure, virtual SANs over the same physical infrastructure. (Please refer to page 94 for more details on this feature.)
- Advanced Diagnostics and Troubleshooting – FC Ping, FC Traceroute, SPAN, RSPAN, and Call Home
- Comprehensive Security - SSH, SFTP, RADIUS, SNMPv3, and Role Based Access Control (RBAC)
- Comprehensive Fabric Management – CLI, SNMP, and Java-based GUI
- Traffic Management – QoS and FCC (Fibre Channel Congestion Control)
- High Availability Software – Failed process restart
- Port-Channel – ISL link aggregation for highly resilient SAN architectures
- SAN Extension – FCIP and iSCSI (MDS 9506, 9509, 9216/9216A/9216i)
- FSPF

All C-Series Directors have the following additional features

- Hitless Software Upgrades – Ability to upgrade without disruption to traffic
- Hot swappable line-cards, supervisors, power supplies, and SFPs
- Redundant supervisor, cross-bar fabric, and power supplies

C-Series Modular Switches also support the following features

- Integrated Multi-protocol capability – FCIP and iSCSI

The following table highlights some of the key differences of the various C-series switch models in terms of features. The differences are mostly hardware and HA (high availability) related. Other features, built in the switch firmware, such as VSANs or Port-Channeling, are common to all models.

Table 6: C-Series switch model features

	Redundant Power, Hot Swap Power	Redundant Cooling, Hot Swap Cooling	Redundant Control Processor	Non-disruptive Code Activation	Port Module Support	Non-blocking	Multi-Protocol Support
MDS 9509 and MDS 9506 Directors	Yes	Yes	Yes	Yes	Yes	Yes	FC, FCIP, iSCSI
MDS 9216/9216A/9216i Fabric	Yes	Yes	No	No	Yes	Yes	FC, FCIP, iSCSI
MDS 9120/40 Fabric	Yes	Yes	No	No	No	Yes	No

M-Series Switch Model Usage

The following two tables show the recommended usage of each M-Series switch model for director or edge switch placement in a fabric. Refer to Table 9 for specific switch model features.

Table 7: M-Series switch model usage as a director switch

M-Series Director Switch Model Selection			
	12-500 Total Ports	501-1000 Total Ports	1000-1632 Total Ports
Director Switch 2/140	Excellent	Excellent	Excellent
Director Switch 2/64	Excellent	Excellent	Excellent
Edge Switch 2/32	Very Good	Good	Good
Edge Switch 2/24	Very Good	Good	Good
Edge Switch 2/12	Good	Not recommended	Not recommended

Table 8: M-Series switch model usage as an edge switch

M-Series Edge Switch Model Selection			
	12-500 Total Ports	501-1000 Total Ports	1000-1632 Total Ports
Director Switch 2/140	Good	Very Good	Excellent
Director Switch 2/64	Good	Very Good	Excellent
Edge Switch 2/32	Excellent	Very Good	Good
Edge Switch 2/24	Excellent	Very Good	Good
Edge Switch 2/12	Very Good	Very Good	Good

M-Series Switch Model Features

All M-Series products come with the following standard features:

- Hot Code Activation Technology (HotCAT™) for non-disruptive activation of new code releases.
- Full non-blocking performance across all ports.
- Consistent latency across all ports
- Redundant and hot swappable power and cooling systems

- Redundant, dual power connections for separate connections
- Hot swappable short and long wave optical transceivers
- Embedded web sever for device and small fabric administration

All M-Series Director Switches have the following features:

- 4-ports per card for low service impact
- Non-blocking port density to minimize data center floor space usage
- Operational requirements with the lowest power consumption and heat generation
- Fully redundant hot swappable switching and processor logic cards
- High availability
- Non-disruptive port expansion
- Non-disruptive failover of redundant components with full performance
- Automatic health checks of redundant field replaceable units (FRUs)

All M-Series Edge Switches have the following features:

- Flexport non-disruptive port expansion
- Edge Switch 2/24 supports Fibre Channel loop connectivity

Table 9: M-Series Switch Model Features

	Min / Max ports	Size 1U =1.75"	Redundant Control Processor, Switching	Non-disruptive Port Expansion	High Availability (Redundant Active Components)
Director Switch 2/140	32-140, 4 port cards	12 U	Yes	Yes	Yes
Director Switch 2/64	33-64, 4 port cards	9 U	Yes	Yes	Yes
Edge Switch 2/32	16, 24, 32	1.5 U	No	Yes	No
Edge Switch 2/24	8, 16, 24	1 U	No	Yes	No
Edge Switch 2/12	4, 8, 12	1 U	No	Yes	No

M-Series Multi-Protocol Support

M-Series switches are supported for iSCSI and FCIP through the SR2122-2 storage router. Refer to Chapter 7, "[SAN Extension](#)" for more information.

Definitions

In order to understand SAN design, it is important to understand the relationship between a SAN and a fabric.

The Storage Networking Industry Association (SNIA) offers the following (slightly reworded) definitions of these two terms. (Refer to www.snia.org)

SAN (Storage Area Network): A network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. The term SAN is usually (but not necessarily) identified with block I/O services rather than file access services. A storage system consisting of storage elements, storage devices, computer systems, and/or appliances, plus all control software, communicating over a network.

Fabric: A Fibre Channel switch or two or more Fibre Channel switches interconnected in such a way that data can be physically transmitted between any two N_Ports on any of the switches. The switches that constitute a Fibre Channel fabric are capable of routing frames using only the D_ID in a FC-2 frame header.

The SNIA definition of "fabric" refers to Fibre Channel and expresses the concept of fabric that is used in this document. An N_Port is an endpoint in the fabric, which is present in a server's Host Bus Adapter (HBA), a storage system, or a storage management appliance connected to the fabric. The D_ID is the address of the destination N_Port, and is contained in the header of every Fibre Channel packet. A Fibre Channel packet is also called a frame. FC-2 is the layer in the Fibre Channel protocol associated with packet routing.

Another way to define a fabric is to view it as a single FC-2 address space. Any valid FC-2 address that can be reached from a given N_Port in a fabric is part of that fabric.

The SNIA definition of "SAN" does not require that a SAN be implemented with Fibre Channel technology. When the term SAN is used in connection with Fibre Channel technology, use of a qualified phrase such as "Fibre Channel SAN" is encouraged. This usage is usually not required, because in most cases today, a SAN is a Fibre Channel SAN.

According to the SNIA definition, an Ethernet-based network whose primary purpose is to provide access to storage elements would be considered a SAN. SANs are also sometimes used for system interconnection in clusters.

A given SAN can contain one or more fabrics. Most small SAN configurations use a single fabric. Larger storage environments may require very high levels of availability, and the best way to obtain this is by using a pair of redundant fabrics. HP storage systems, servers, and operating systems support both design approaches.

HP Standard SAN Topologies

The HP standard topology designs reflect the proper application of the HP SAN design rules. Each of the standard designs is tailored for a particular data access and connectivity need. Collectively, these designs provide a wide range of options for selecting the appropriate SAN design for your specific requirements. Variations of these designs, including additions or changes, can be validated by adhering to the appropriate rule set for each topology type. A subset of a standard design is always acceptable.

The different types of HP standard SAN topologies are described in detail in the following sections.

SAN Fabric Topologies

SAN fabric topology designs include:

- Cascaded Fabrics
- Meshed Fabrics
- Ring Fabrics
- Backbone Fabrics

Each of the design types can be:

- Implemented as a separate SAN for specific departments or applications within a company, to accommodate different data access needs.
- Implemented with centralized backup capabilities, reducing the cost of backup and restore operations.
- Deployed in one or more co-located groups.
- Connected to another SAN using routing technology. (B-Series fabrics)
- Deployed in multiple VSANs. (C-Series fabrics)
- Deployed across a wide area with extended inter-switch distances

Note: Refer to Chapter 7, "[SAN Extension](#)", for additional information on extending SANs over long distances.

- Used to begin an ongoing deployment process using SANs and Fibre Channel technology in a modular, controlled approach. Storage consolidation can be implemented on a departmental or independent SAN basis. Future capabilities will allow for more switches within a single SAN, interconnection of multiple SANs to build larger fabrics, and provide for additional consolidation, if desired, or broader server-to-storage access.
- Centrally managed.
- Implemented with a wide range of SAN availability levels. See "[Levels of Availability](#)."
- Upgraded to higher capacity topologies or topologies optimized for different data access types if needs change.

Single-Switch Fabrics

The smallest SAN consists of a single Fibre Channel switch, server, and storage system. This topology is a subset of all the other topologies, and forms the basis of the range of HP SAN solutions.

By choosing among various HP Fibre Channel switches, you can construct a wide range of single-switch SAN solutions. The smallest supported HP SAN uses a single 8-port switch, the HP StorageWorks SAN switch 2/8. If you need a large single switch SAN, the HP StorageWorks SAN Director 2/140 offers 140 ports. Furthermore, if there is a need for high SAN availability, two independent single-switch SAN fabrics may be used in a dual-fabric environment to give a total of 280 user ports.

A single-switch fabric maximizes SAN performance, because every port on the switch has full connectivity to every other port on the switch. This design is also very easy to install and configure, since there are no connections from one switch to another.

An example of a simple single-switch SAN is shown in Figure 1.

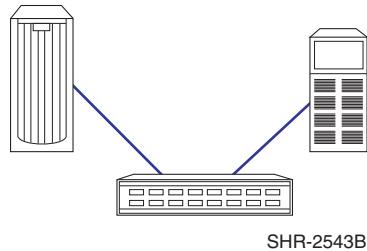


Figure 1: Single-switch SAN

Starting from a single-switch configuration, you can add more switches to your SAN fabric—following the support limits listed for each of the fabric topology designs—to increase the number of connections for servers and storage.

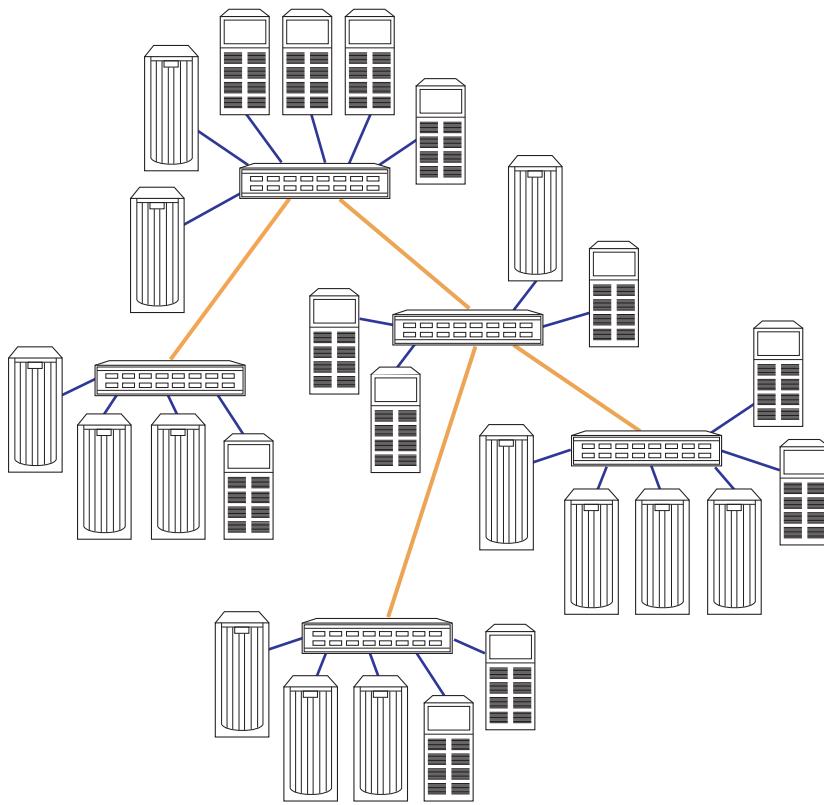
You can also view the various topology options as a way to connect existing smaller SANs or SAN islands. If you have already deployed several small SANs, you can connect them together to make a larger SAN. For example, if you have two single-switch SANs, you can connect them together into a cascaded fabric. Or, if you have deployed two four-switch meshed SANs as separate SANs, you can merge these into a larger single 8-switch meshed SAN as shown in Figure 4, "Modified Meshed Fabric SAN". If you have multiple single-switch SAN fabrics, you can connect these into a single larger SAN fabric by connecting them in a ring, or to a central backbone, using the backbone fabric topology.

Cascaded, Meshed, and Ring SAN Fabrics

The first three fabric topologies that involve more than a single switch are organized so that all of the switches in the fabric are used for connecting servers and storage. Every switch has at least one user port. Typically, in these types of fabric arrangements, a small percentage of the total number of switch ports is used for inter-switch connectivity in the form of Inter-Switch Links (ISLs). Refer to Chapter 1, "Connectivity", for more information about trading ports used for ISLs for ports used for servers and storage.

Cascaded Fabrics

A cascaded fabric SAN (see Figure 2) is a set of switches connected together, by one or more ISLs, in a tree-like arrangement.



SHR-2552B

Figure 2: Cascaded Fabric SAN

Cascaded fabric designs are well suited to environments with local data access patterns. In these cases, I/O requests from a servers attached to a given switch are made most often to storage systems that are attached to the same switch. Groups of servers and their storage systems can be connected to the same switch to provide the highest level of I/O performance. Cascading provides a means to scale the SAN for additional connectivity of servers and storage, and allows for centralized management and backup, while maintaining the high I/O performance of local access.

Cascaded designs can also be used for centralized or distributed access; however, traffic patterns should be well understood and should be factored into the design to ensure that there are an adequate number of ISLs to meet performance requirements. Using more than one ISL between switches in a cascade also provides redundant paths between a given pair of switches in the fabric. HP highly recommends that cascaded designs be implemented with a minimum of two ISL connections on each switch, either as a pair of ISLs between the same two switches or by connecting every switch to at least two other switches in the fabric.

Very Large Cascaded Director Plus Edge Switch Fabric

The largest fabric supported by HP at this time is based on a cascaded fabric using HP StorageWorks Director 2/140 and StorageWorks Edge switches. Because the cascade configuration may be used to maximize user port count—if relatively low bandwidth between user ports that are on separate Directors is acceptable—this configuration can provide a level of connectivity that cannot be obtained by other means. Because of the availability features built into the SAN Director product, this configuration also offers high availability in a single-fabric topology.

The SAN Director product family is limited to a three-hop maximum. Hop count is a measure of distance or links between switches, measured in ISLs. A cascade configuration with 24 SAN Directors and Edge switches that meets the three-hop rule has one Director at the top of the tree, with 7 Directors attached under it, and 16 Edge switches under them. This gives a total of 24 SAN switches in one fabric. This supported configuration provides 1024 user ports.

Advantages of Cascaded Fabrics

- Accommodates diverse geographic conditions
- Scales easily for additional connectivity
- Shared backup is supported
- Shared management is supported
- Optimal local access is inherent in the fabric design
- Most efficient in cost per port

Meshed Fabrics

In a meshed fabric design, all of the switches are interconnected so there are at least two paths or routes from any one switch to any other switch in the fabric. This type of connectivity provides fabric resiliency. If a single ISL or ISL port interface fails, the fabric can automatically re-route data through an alternate path. The new route can even pass through additional switches in the fabric. An example of a meshed fabric is shown in Figure 3.

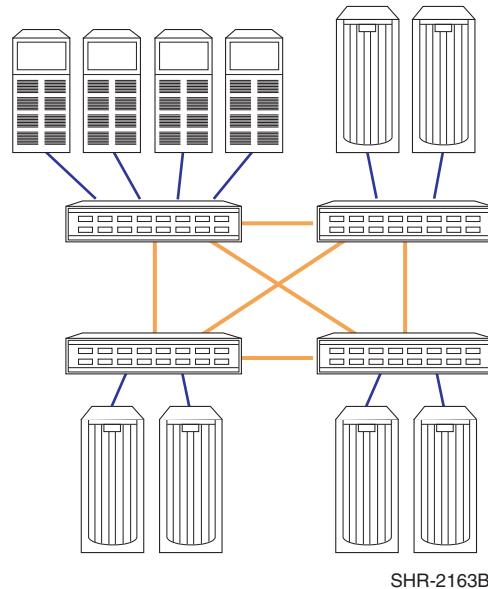


Figure 3: Meshed Fabric

As switches are added to a meshed topology, the number of ISLs required to maintain full connectivity between any switch and any other switch becomes excessive. This reduces the number of user ports in comparison to the total number of ports, which is a measure of the connection efficiency of the fabric. The connection efficiency of this fabric design can be improved by implementing a slightly modified mesh design, as shown in Figure 4. In this case the connectivity between switches is reduced, but the fabric availability is maintained because there are still multiple paths between switches.

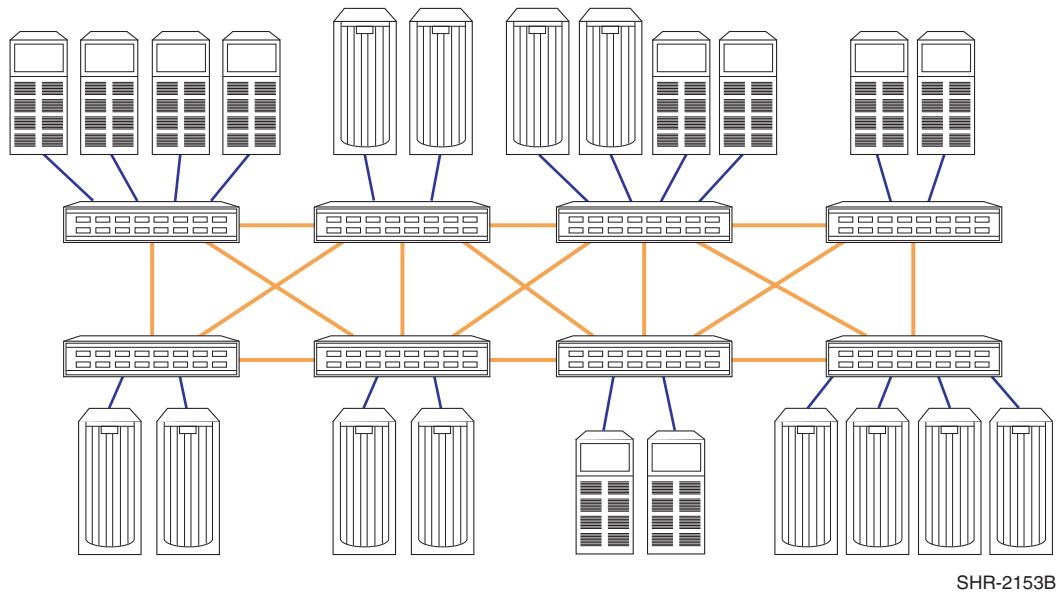


Figure 4: Modified Meshed Fabric SAN

In this example diagram, as switches are added, they are only connected to adjacent switches, not all other switches in the fabric. This still provides the benefits of full many-to-many connectivity without a decrease in connection efficiency.

Meshed fabrics are well suited to applications where data access is a mix of local and distributed. The full connectivity (or high connectivity, in the case of modified meshes) supports many-to-many access, while at the same time allowing localized access to individual switches, servers and storage.

Advantages of Meshed Fabrics

- Can be configured for any to any or local data access, or a mix
Reduces staff effort by minimizing reconfiguration and re-cabling of existing Fibre Channel switches. Adapts easily to new or different storage needs.
- Provides protection against link and switch port failures
Fabric design allows Fibre Channel switches to automatically re-route under failure conditions, saving time and effort to manually trace the problem and re-route.
- Scales easily
The mesh design can be extended from a four-switch fabric to six or eight switches easily, and without disruption to the existing SAN. The mesh design affords ease of adding servers to the SAN without impacting existing connections or equipment. This is especially useful for companies where there is rapid growth, or computing and storage needs are changing frequently.
- Shared backup is supported
One or more Automated Tape Libraries can be added to the mesh fabric at various points without impacting performance or management.
- Shared management is supported
All Storage Management Appliance tools can navigate and manage the Storage Area Network in the mesh fabric, saving time and effort.
- Optimal distributed access is inherent in the fabric design

Ring Fabric

A ring fabric (see Figure 5, "Ring Fabric SAN") is a continuous ring of switches connected together into a single fabric. Each switch is connected to adjacent switches, with the last switch in the ring connected back to the first. This arrangement of switches provides almost the same level of fabric resiliency as the mesh design, with full fabric connectivity and at least two internal fabric paths or routes.

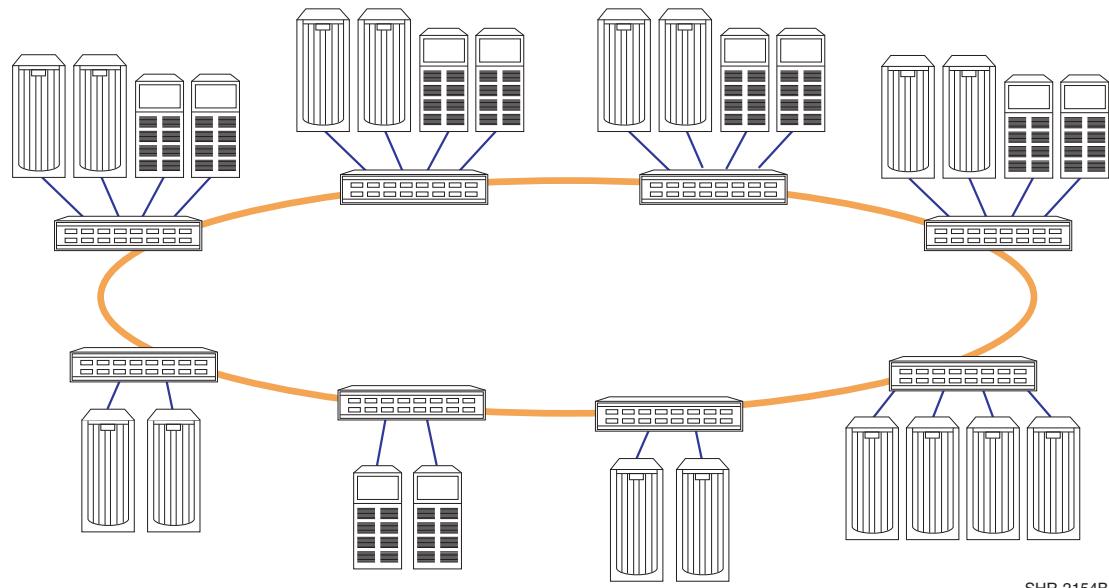


Figure 5: Ring Fabric SAN

If you use fewer than 12-switches in a ring constructed using B-Series product line switches, you can add additional switches to the outside of the ring. These satellite switches provide additional user ports with only a slight reduction in fabric availability. For example, 11 satellite switches can be connected to a 11-switch ring. This results in a 22-switch fabric and maintains the overall seven hop limit. Figure 6 shows a 22 -switch fabric.

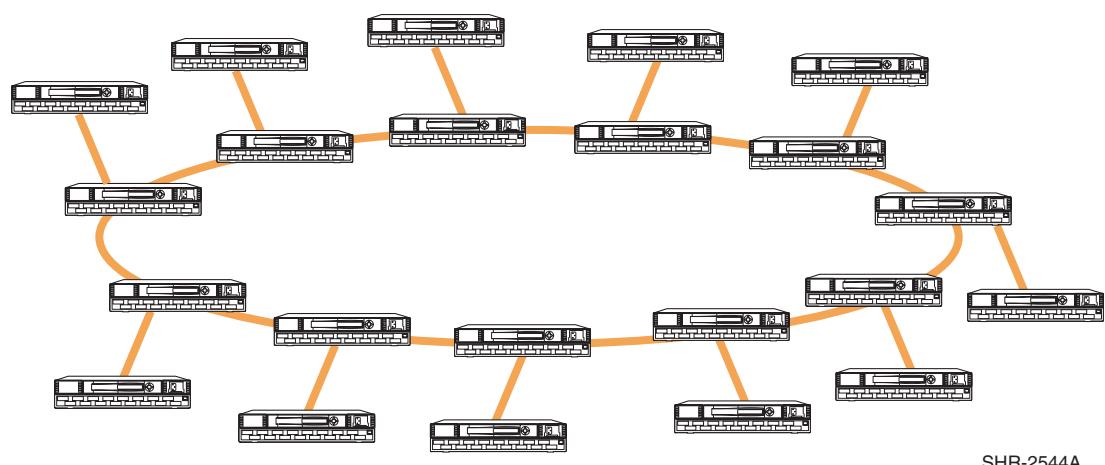


Figure 6: Ring Fabric SAN with Satellite Switches

Ring fabric designs are well suited to applications where data access is always localized. Servers and the storage that is accessed are on the same switch, and the majority of data traffic is handled within that switch. This implementation provides a way to scale the fabric in a modular fashion by adding a switch and groups of servers and storage as a cell, using a building block approach to increase the size of the SAN over time. This is particularly useful in situations where the storage capacity requirements vary over time, such as in a storage service provider environment.

A ring fabric can be pre-configured and installed before the server requirements are known. This is useful because the ability to install the fabric infrastructure beforehand can greatly simplify the installation of each incremental storage system or server. Interconnecting the switches in a ring topology provides a communication path that supports centralized SAN management and centralized backup.

The ring fabric is not recommended for applications that require many-to-many connectivity.

Advantages of Ring Fabrics

- Easy to build
 - Each Fibre Channel switch can support servers and storage, thus saving time and effort on SAN design and implementation.
- Scaling is simple and non-disruptive
 - Fibre Channel switches can be added one at a time, as storage and connection needs dictate. Each Fibre Channel switch can support identical servers and storage for controlled growth, or can support a variety of heterogeneous systems for new demands of the business.
- Shared backup is supported
 - One or more Automated Tape Libraries can be added to the ring fabric at various points without impacting performance or management.
- Shared management is supported
 - All Storage Management Appliance tools can navigate and manage a SAN with a ring topology, saving time and effort.
- Optimal local access is inherent in the fabric design
 - The majority of the data traffic is within each switch in the ring, minimizing any allocation, fabric and performance issues.
- Modular design
 - Saves time and effort on design and implementation by complementing the basic modularity of all StorageWorks products, including the raid array controllers, universal packaging, and secondary storage (Automated Tape Libraries).

Backbone Fabric

A backbone fabric has one or more Fibre Channel switches primarily dedicated to connecting to other switches within the fabric. The backbone switches provide high bandwidth and redundant connectivity to the other switches. This type of implementation offers the best "many-to-many" connectivity.

Backbone fabrics are well suited for implementations where the primary requirement is for full network “many-to-many” connectivity with high performance. They are the most conservative design approach in cases where the I/O traffic patterns are unknown or varying. They are also the best design to choose if you plan to implement SAN-wide storage pooling and sharing, and for environments that use storage virtualization.

Figure 7 shows a backbone fabric where the switches in the center are dedicated to providing connections between the other switches. The switches to which servers and storage can be connected are called “edge switches”, and the switches in the center are called “backbone switches”. Servers and storage can be connected to the user ports on any of the edge switches, which maximizes the flexibility of how you use the user ports.

If required, you may choose to connect centralized primary (disk) or secondary (tape) storage directly on the backbone switches. This approach may be useful if excess ports are available on the backbone switches.

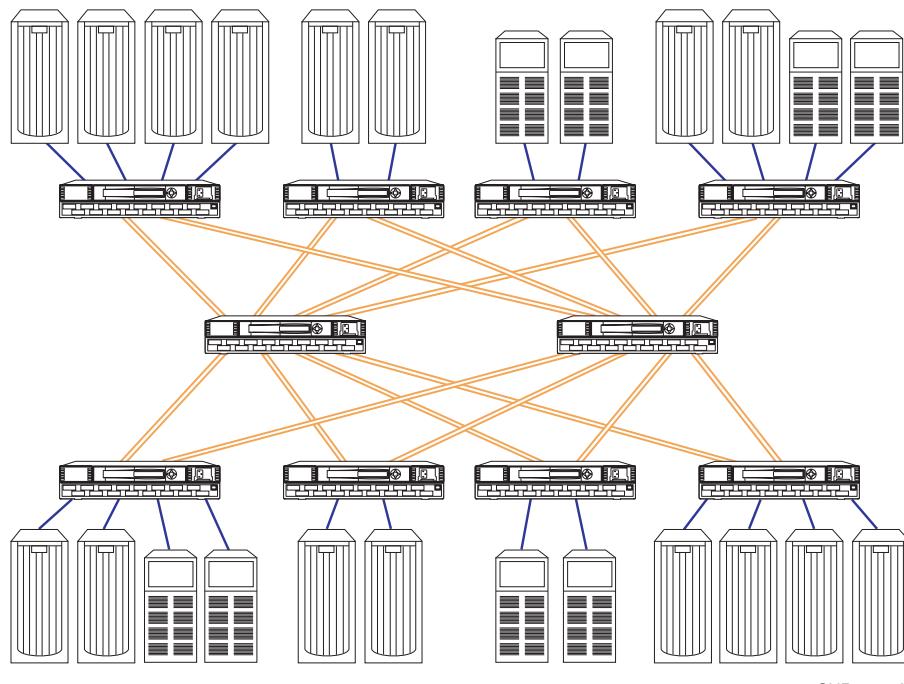


Figure 7: Backbone Fabric SAN

HP StorageWorks SAN fabrics currently support a large number of switches in a fabric. Figure 8 shows a large backbone fabric SAN with 20 switches. Configurations of this type can be used to support the most demanding requirements for storage system size and performance.

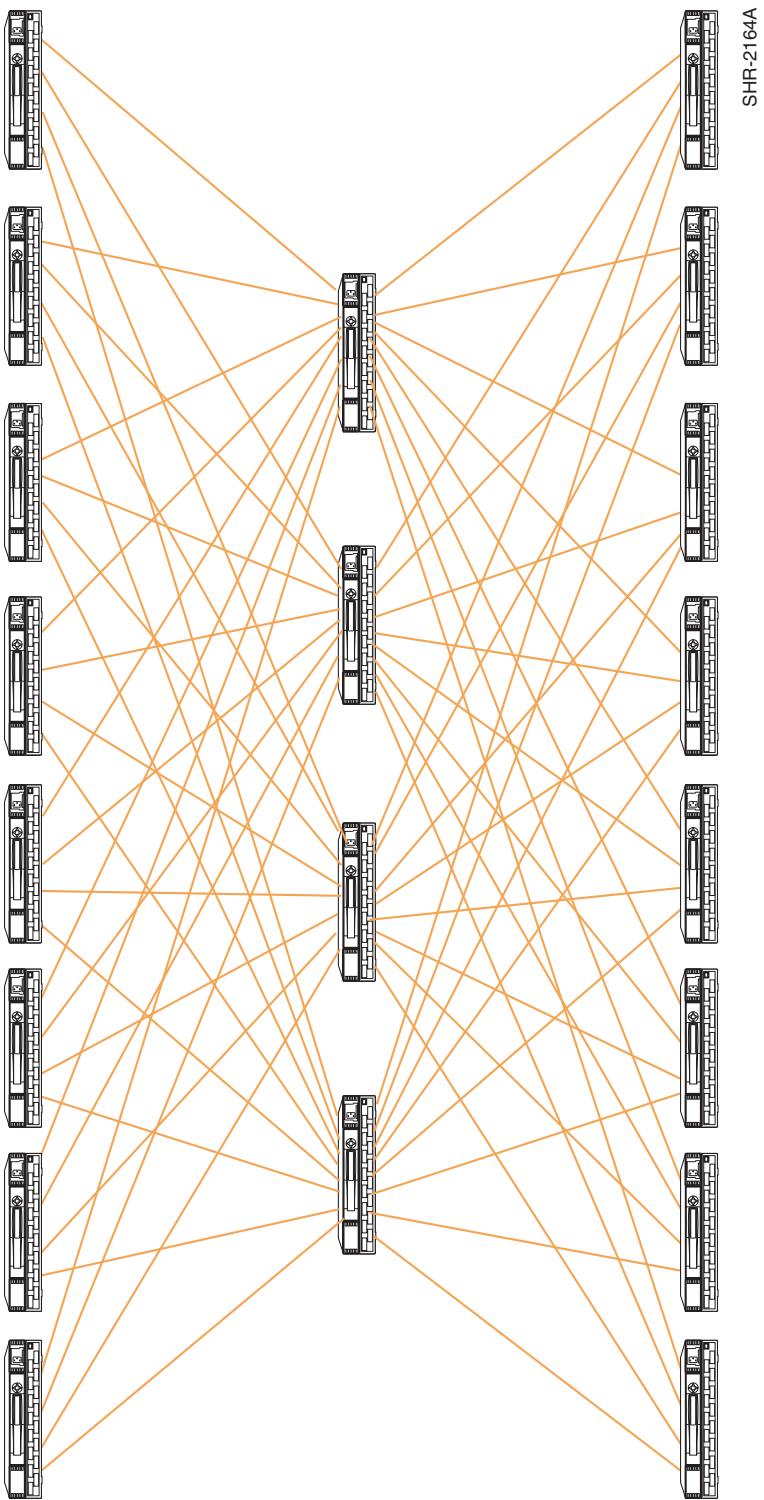


Figure 8: Backbone SAN with 20 Switches

Fat Tree and Skinny Tree Designs

The SAN shown in Figure 7 can be drawn with the backbone switches at the top and the edge switches collected together into a row at the bottom, as shown in Figure 9. This method of illustrating the topology is helpful when evaluating the potential performance of a given backbone configuration. It shows the switches in a hierarchy, where the edge switches form a layer that provides access to the SAN and the backbone switches form a layer that distributes I/O requests between edge switches.

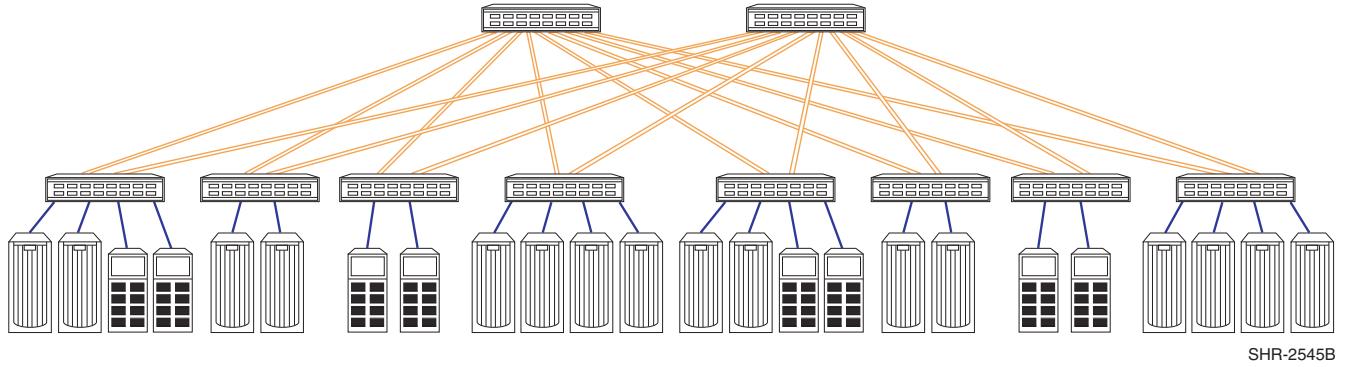


Figure 9: Backbone SAN, Drawn Hierarchically

Depending on how a backbone SAN is designed it can be classified as a “fat tree” or a “skinny tree”. The difference between fat and skinny trees is the number of ISLs used to connect the edge switches to the backbone switches. The number of ISLs subtracts from the number of end ports and therefore affects the total number of switches needed for a particular configuration. Fat trees use 50% of the edge switch ports as ISL connections while skinny trees use fewer than 50%.

Figure 10 shows two hierarchical SAN fabrics. The skinny tree illustration shows that if all the devices connected to user ports on the left-hand side of the tree want to communicate with the devices on the right-hand side of the tree, then there are not enough connections to the switch at the top of the tree for the required traffic. There are too many port-pairs trying to use the ISLs on the switch at the root of the tree. The fat tree illustration shows that by providing additional ISLs between the switches that are further up in the hierarchy, full-performance bandwidth can be provided for *any combination of port-pairs*. This important benefit is the basis for the use of fat tree configurations in application environments where the highest level of performance and capability are required.

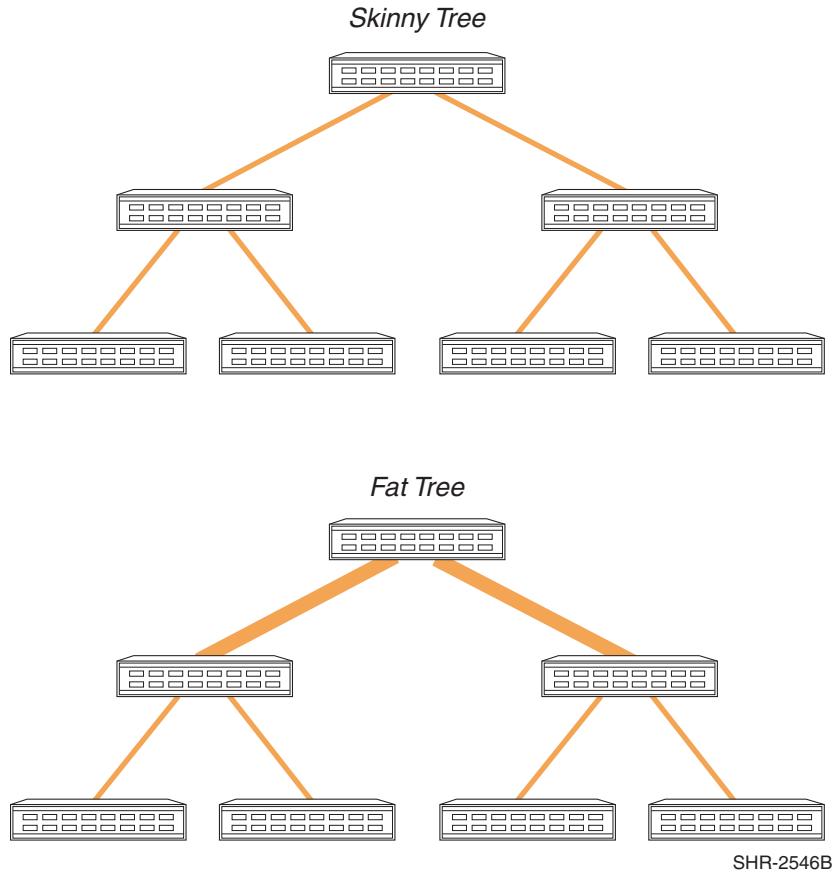


Figure 10: Skinny Tree and Fat Tree

This distinction in the number of ISL connections between fat and skinny trees results in two major differences:

1. Skinny trees require fewer switches than fat trees to supply the same number of user ports. Figure 11 shows how six 16-port switches in a skinny tree configuration yield 64 user ports, while the same switches wired in a fat tree as shown in Figure 12 yield only 32 user ports.
2. Fat trees have more ISL connections and therefore have higher cross sectional bandwidth capabilities than skinny trees. The term cross sectional bandwidth is used to refer to the maximum amount of data that can pass through the ISL connections at the midpoint of the fabric.

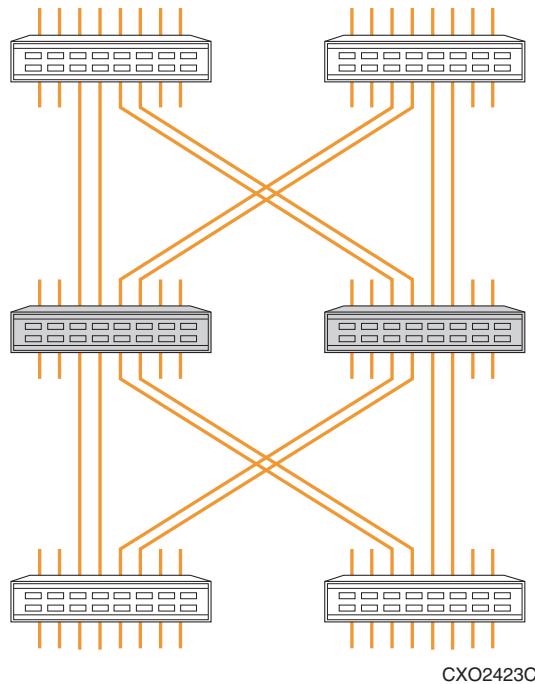


Figure 11: 64-Port Skinny Tree

The configuration shown in Figure 11 has six 16-port switches. Two are backbone switches (shaded) and four are edge switches. The edge switches each have 12 user ports available to connect to servers or storage. With the 48 ports on the edge switches and with the 16 available ports on the backbone switches the total number of user ports is 64.

There are 8 ISLs on each side of the backbone switches.

Note that a mix of 1 Gbps and 2 Gbps switches or devices makes the evaluation of this type of configuration more complicated.

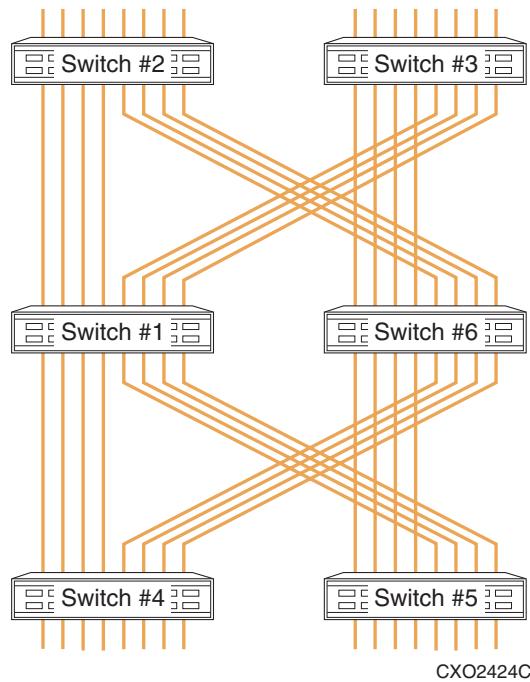


Figure 12: 32-Port Fat Tree

The configuration shown in Figure 12 also has six 16-port switches, with two backbone switches and four edge switches. The edge switches each have eight user ports available to connect to servers or storage. No ports are available on the backbone switches for user ports, so the total number of user ports is 32.

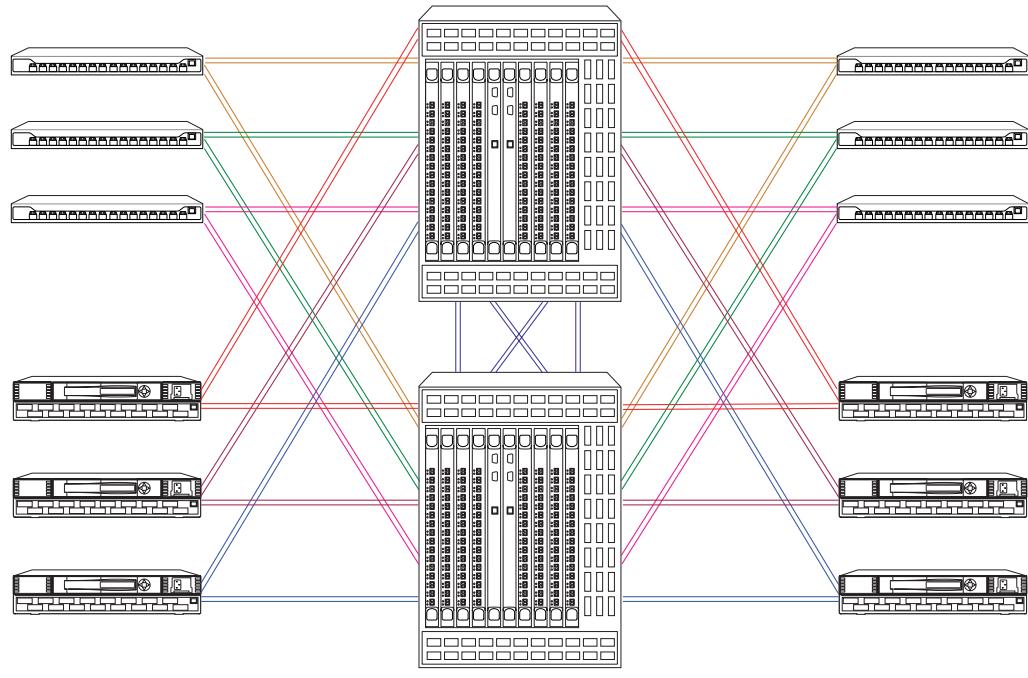
The two configuration designs shown are utilized in the legacy HP Surestore FC Switch 6164 and the legacy Compaq StorageWorks SAN Switch Integrated/32 and 64.

Backbone SANs Using Core Switches and Directors

With the introduction of switches with a higher number of ports, backbone fabrics may be constructed using a mix of small and large switches. The HP standard SAN designs using this approach are as follows.

The “4 x 12” configuration, shown in Figure 13, uses a mixture of HP StorageWorks Core Switch 2/64 and SAN Switch 2/16 or SAN switch 16 devices. Four core switches are connected in a mesh configuration with two ISLs between each pair of switches in the mesh. Twelve SAN switches are connected to the mesh with four ISLs between each SAN switch and the mesh. This gives a total of 304 user ports, including 160 on the mesh and 144 on the SAN switches.

This configuration is useful for situations where a large port count is required in a single fabric, and where many high-performance systems can make use of direct connections to the mesh. For example, high performance storage systems can make good use of a 2 Gbps connection, and should be connected directly to the mesh. Since each 16-port SAN switch has four ISLs connecting it to the core, the 8 Gbps (4 x 2 Gbps) total bandwidth between the SAN switch and the mesh must be shared between the 12 user ports on that switch. This means that the I/O performance requirement for each server attached to a SAN switch in this configuration must be reviewed to make sure that the ISLs are not an I/O bottleneck. Because of the limited number of ISLs between the switches in the mesh, this configuration is not appropriate for environments where many-to-many traffic patterns are predominant.



SHR-2554A

Figure 13: “4 x 12” Backbone SAN

The “4 x 24” configuration also uses a mixture of HP StorageWorks core switch and SAN switch devices. By using a larger number of switches, this topology design provides adequate internal connectivity within the mesh and also provides more user ports than the 4 x 12 configuration.

In the 4 x 24 configuration, shown in Figure 14, four core switches are connected in a mesh with four ISLs between each pair of switches in the mesh. 24 SAN switches (1 Gbps or 2 Gbps models) are connected to the mesh, with a total of four ISLs on each SAN switch (just the same as in the 4 x 12 configuration). This gives a total of 400 user ports, including 112 on the mesh and 288 on the SAN switches. This configuration may provide a better trade-off of high-performance connections—directly to the core switches—and lower-performance connections on the SAN switches.

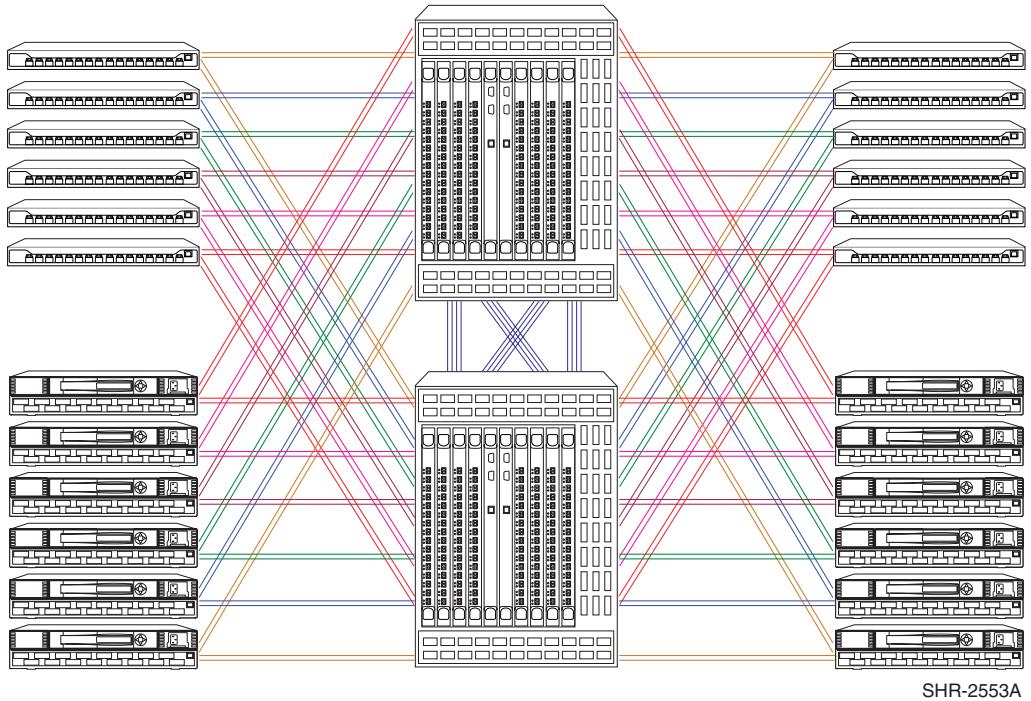


Figure 14: "4 x 24" Backbone SAN

Director Fabrics

Backbone fabrics may be constructed using Director and Edge switches. These products may be used in configurations where no more than half of the ports on a given switch is used for ISL connections. This means that a fat tree configuration cannot be constructed because the distribution layer cannot be fully populated. On the other hand, there are fewer topology restrictions for these products, so you can combine a mixture of Director and Edge switches, with a single Director as the backbone, to make a skinny tree fabric.

For example, a skinny tree using the design shown in Figure 11, but using 64-port Directors, meets the ISL port count and hop count requirements for this family of products while providing 256 user ports.

Figure 15 shows a "Director plus Edge switch" tree design using two HP StorageWorks Directors and 14 HP StorageWorks Edge switches. In this topology, high performance storage systems are connected directly to the Directors, because they can make use of the full bandwidth of the 2 Gbps Fibre Channel connections. Servers are connected to the edge switches, because they require only a smaller amount of bandwidth. Depending on the bandwidth requirements of the servers, the number of ISLs between the edge switches and the Directors can be varied.

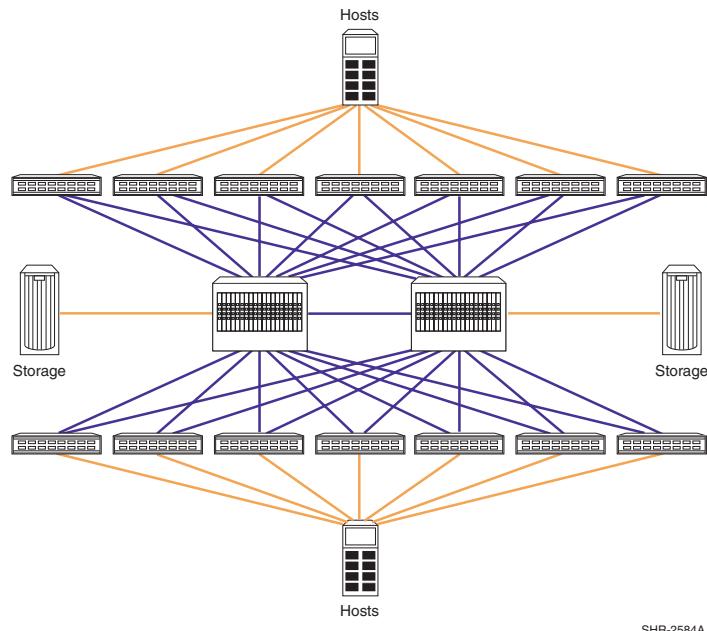


Figure 15: Director plus edge switch SAN

Advantages of Backbone SANs

- Efficient port expansion: new switches need only be connected to backbone switches.
Saves time and effort during the design and implementation phases by isolating the new switches from the existing SAN backbone.
- All edge switches are only two hops apart.
Saves design effort for adding new servers and storage to any point on the SAN. The uniformity of access supports new usage patterns without requiring redesign and re-cabling.
- When implemented with two or more backbone switches, provides a level of switch redundancy in a single fabric.
Backbone design allows Fibre Channel switches to automatically re-route under failure conditions, saving time and effort to manually trace the problem and re-route.
- Can be centrally managed.
All Storage Management Appliance tools can navigate and manage the Storage Area Network in a tree backbone fabric, saving time and effort.
- Full "many-to-many" connectivity with evenly distributed bandwidth and redundant connectivity.
Supports varying connection and performance demands regardless of the location within the SAN. At the same time, provides uniform routing and redundancy from a single SAN design.
- Improved bandwidth with multiple parallel ISLs.
Additional ISLs ensure that all data traffic within the tree backbone SAN will be managed with less performance degradation, regardless of the location of servers and storage relative to each other.

- Offer maximum flexibility for implementing mixed access types: Local, Distributed, or Centralized.
Saves effort planning data traffic patterns; the tree backbone supports all access patterns.
- Can be implemented with centralized backup capabilities, reducing the cost of backup and restore operations
- Can be implemented with all availability levels
Saves effort in the design and implementation phases by offering a single design for a variety of usage requirements.
- Can be an upgrade path from other SAN designs. Backbone SAN designs offer evenly distributed bandwidth and full many-to-many connectivity; they are the best solution for flexible SAN-wide storage pooling and sharing.
- Well-suited to take full advantage of expected future technological developments such as storage virtualization
Saves the investment made in the SAN by continuing its use as more advanced tools, products, and designs become available.

Topology Data Access Usage

The various SAN topology options can be characterized by how well they support specific data access patterns. Refer to Chapter 1, "[Data Locality](#)". Table 10 provides a general characterization of the different topology designs as a means to compare each of the design types by optimal data access capabilities. Use the table as a basis for selecting the best-suited topology for your expected access needs.

Individual topologies can be tailored or modified to better meet specific requirements. For example, choosing a fat tree backbone design provides the best overall “many-to-many” connectivity, and allows portions of the tree implementation to be configured for local access. This can be accomplished by connecting servers and storage typically accessed on the same switch within portions of the tree backbone.

Table 10: Topology Usage Rating

SAN Topology	Data Locality		
	Local “One-to-One”	Centralized “Many-to-One”	Distributed “Many-to-Many”
Cascaded	Highest	Not Recommended	Not Recommended
Meshered	Medium	Medium	High
Ring	Highest	Medium	Not Recommended
Skinny Tree Backbone	Medium	High	High
Fat Tree Backbone	High	Highest	Highest
Single Switch	Highest	Highest	Highest

Topology Maximums

Table 11, Table 12, and Table 13 indicates the maximum number of switches and ports supported for each of the HP standard SAN topologies.

Note: The maximums shown assume the use of the minimum number of ISLs. Depending on your specific application, you may need more ISLs. This reduces the overall number of ports available for servers and storage. Attaching the Storage Management Appliance also reduces the total number of ports available for servers and storage. See Chapter 5, "SAN Management."

Table 11: Topology Maximums when using B-Series Product Line Switches

SAN Topology	Maximum Number of Switches	Maximum Total Number of Ports	Maximum Number of User Ports
Single Switch	1	128	128
Cascade	34	1280	1262*
Mesh	34	1280	1238*
Ring	15	1280	1260*
Ring with Satellite Switches	34	1280	1260*
Backbone (4 by 30 with 4 core/director switches and 30 SAN switches)	34	1280 (Refer to Chapter 3, SAN Fabric Design Rules, for specific configuration rules.)	1160

Note: *While this is a valid configuration, it achieves a high user port count by severely limiting the connectivity within the SAN by using 1 ISL between each switch pair.

Note: Refer to Chapter 3, SAN Fabric Design Rules, for specific configuration maximums.

Table 12: Topology Maximums when using M-Series Product Line

SAN Topology	Maximum Number of Switches	Maximum Total Number of Ports	Maximum Number of User Ports
Single Switch	1	140	140
Cascade	24 (maximum of 8 Directors)	1,632	1024 (cascade with 8 Directors, plus 16 edge switches)
Mesh	Not Applicable, Exceeds Hop Count Limit	Not Applicable	Not Applicable
Ring	7	980	966
Ring with Satellite Switches	Not Applicable	Not Applicable	Not Applicable
Backbone	24	1632	1024

Note: Refer to Chapter 3, SAN Fabric Design Rules, for specific configuration maximums.

Table 13: Topology Maximums when using C-Series Product Line Switches

SAN Topology	Maximum Number of Switches	Maximum Total Number of Ports	Maximum Number of User Ports
Single Switch	1	224	224
Cascade	24	3000 (maximum of 12 Directors)	2500 (cascade with 12 Directors and 10 fabric switches)
Backbone	24	3000 (maximum of 12 Directors)	2500

Data Availability in a SAN

Data availability in a computer installation is influenced by many factors, including the application software and operating systems in the servers, the server hardware, the SAN fabric infrastructure, and the primary and secondary storage. Operational parameters including backup schedule and machine room procedures, as well as personnel issues and overall administrative practice all make important contributions to the availability of data in a computer system environment.

In some environments, adequate data availability is established by a routine backup procedure performed on a scheduled basis. In other cases, online dynamic backup of primary data to a remote site is required. Some environments use clustered servers and redundant fabrics in their SAN systems in order to achieve their data availability goals.

When considering SAN fabric topology selection, the number of Fibre Channel switches and the number of ISLs between the switches have the largest effect on the data availability. The number of connections or paths between a given server or clustered servers and the fabric, and the number of storage controller connections or paths into the fabric also affect data availability.

From the perspective of SAN architecture and fabric topology design, fabric availability can be classified into at least four categories or levels. The different categories offer a range of availability levels from the most basic interconnect scheme with no redundancy, up to fully redundant No Single Point Of Failure (NSPOF) designs.

Note: Refer to "[Router High Availability Configurations](#)", page 77 for information on high availability configurations when using the HP MP Router.

Levels of Availability

1. Single Fabric/Single Server and Storage Paths
2. Single Meshed Fabric/Single Server and Storage Paths
3. Single Meshed Fabric/Multiple Server and Storage Paths
4. Multiple Fabrics/Multiple Server and Storage Paths

Level 1: Single Non-meshed Fabric/Single Server and Storage Paths

These designs are implemented with single links between each switch, connected in one fabric. The Fibre Channel switches are arranged so that servers and storage connect into the fabric using single paths. This type of design does not provide any level of fabric or fabric path redundancy.

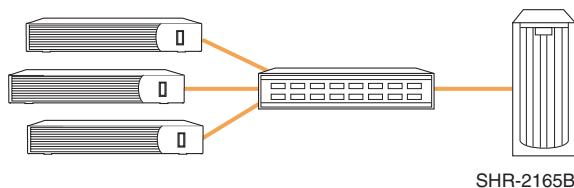


Figure 16: Level 1: Maximum Connectivity

Level 2: Single Meshed or Cascaded Fabric/Single Server and Storage Paths

These designs have more than one ISL between switches and/or multiple paths or routes to all switches in the fabric. Servers and storage connect into the fabric using single paths. This provides the benefit of fabric resiliency. If a single switch port or a link between two switches fails, the fabric automatically re-routes data to an alternate fabric link or route. The servers see no interruption in their I/O flow.

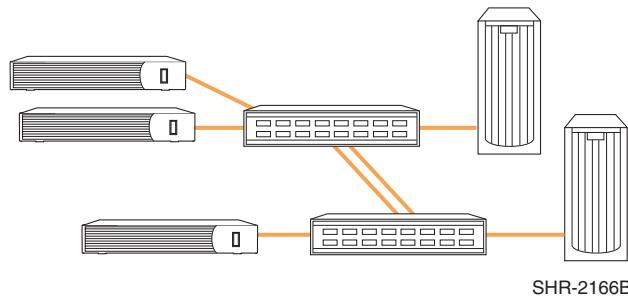


Figure 17: Level 2: Fabric Resiliency

Level 3: Single Meshed or Cascaded Fabric/Multiple Server and Storage Paths

These designs are the same as Level 2 with the addition of multiple data paths between servers and storage connecting into one fabric. Level 3 offers the benefits of both fabric resiliency and multiple server and storage paths. In the unlikely event of a switch, host bus adapter, or path failure, data is automatically re-routed to an alternate path in the servers and storage, and through the fabric. The servers see no interruption in their I/O flow. Level 3 may require (depending on the O/S) the use of fabric zoning to define a minimum of two separate zoned paths for each server configured with multiple paths in a single fabric. To ensure high availability, each HBA must be cabled to a different switch and be configured for access to a different storage system controller when set in multiple-bus failover mode. Each controller must be cabled to a different switch, as shown in Figure 18.

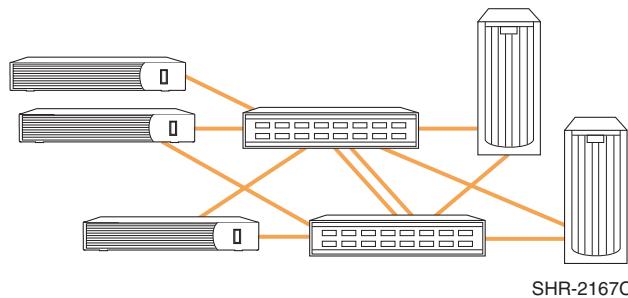


Figure 18: Level 3: Single Fabric High Availability Multi-Pathing

Level 4: Multiple Fabrics/Multiple Server and Storage Paths

Like Level 3, Level 4 provides for multiple data paths between servers and storage, but in the Level 4 designs these paths are connected to physically separate fabrics. This type of design provides the highest level of availability and offers no single point of failure protection (NSPOF). Any event that may affect the fabric performance or usability will be overcome by routing data to another alternate fabric. The servers see no interruption in their I/O flow.

The Level 4 design eliminates any vulnerability to fabric failures, for example, human error such as improper switch replacement procedure, inadvertent erroneous fabric configuration settings, or a fabric service failure. This type of design also provides the highest level of performance and a higher number of available ports, since all fabrics can be accessed and utilized simultaneously during normal operations. This also allows for nondisruptive upgrades.

This level of protection is available for all HP standard SAN topologies by replicating the chosen design in two separate fabrics. HP recommends that the two fabrics have similar or identical topologies. While the cost of the implementation may increase, the added benefit beyond the increase in data availability is an increase in total available ports. For example, choosing to implement a single meshed fabric design using four switches provides up to 52 ports for server and storage connectivity. Implementing the same topology using two fabrics provides up to 104 ports for server and storage connectivity. This is shown in Figure 19.

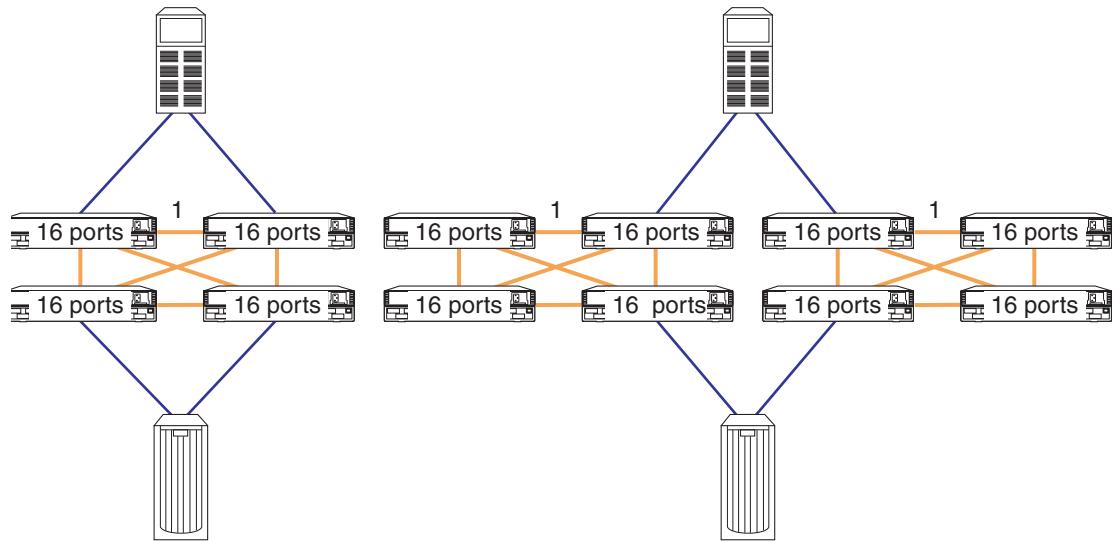


Figure 19: Single Fabric and Dual Fabric SANs

Using two fabrics allows for non-disruptive software and firmware code updates. For example, given the two fabrics shown in Figure 20, you can failover operations to Fabric B, upgrade Fabric A, then fallback operations to Fabric A. The procedure can then be repeated in reverse to upgrade Fabric B.

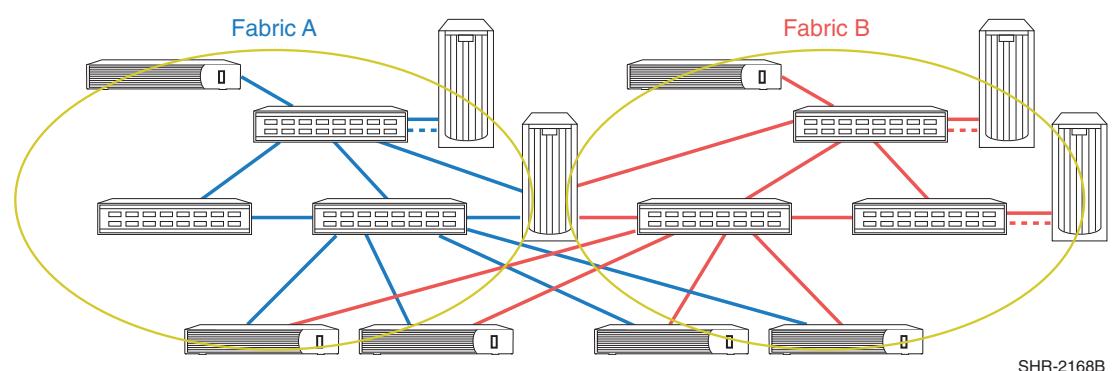


Figure 20: Level 4: Dual Fabric High Availability Multi-Pathing Fault Tolerant

See also Chapter 3, [Figure 34](#).

Table 14 characterizes data availability and indicates the supported topologies for each level.

Table 14: Fabric Design Data Availability

Fabric Design	Availability Level	SAN Topologies
Single Fabric (Non-Meshed)	1	No Redundancy
Single Meshed Fabric Multiple Fabric Paths	2	Medium
Single Meshed Fabric Multiple Fabric Paths Multiple Server and Storage Paths¹	3	High
Two (or more) Fabrics Multiple Server and Storage Paths	4	Highest (NSPOF)

1. May require the use of zoning to define a minimum of two separate data paths within the single fabric. This is platform dependent.

Availability Design Considerations

Two major considerations in choosing an availability level are the criticality of data access and cost. For mission critical applications, first consider full redundant fabric designs. The additional cost can usually be justified when you consider the cost associated with the loss of access to critical data.

You should also remember that the additional cost of more than one fabric provides more than redundancy since the number of available ports will typically double. If this increased connectivity can be utilized by adding more servers and storage to the SAN, the cost factor is minimized. Figure 15 characterizes data availability levels relative to cost and total number of available ports.

Table 15: Availability Cost Factors

Fabric Design	Level	Hardware Cost Factor ¹	Available Ports ²
Single Fabric (Non-Meshed)	1	x	n - #ISL Ports
Single Meshed Fabric Multiple Fabric Paths	2	x + Additional ISLs	n - #ISL Ports
Single Meshed Fabric Multiple Fabric Paths Multiple Server and Storage Paths³	3	x + Additional ISLs + Additional HBAs	n - #ISL Ports - Additional HBA Ports
Two (or more) Fabrics Multiple Server and Storage Paths	4	x + Additional ISLs + Additional HBAs + Additional Switches	2n - #ISL Ports - Additional HBA Ports

1. The variable x is the cost of a single non-meshed fabric. It is used as a reference for comparison.
 2. The variable n is the total number of ports available for devices in a SAN fabric.
 3. May require the use of zoning to define a minimum of two separate data paths within the single fabric. This is platform dependent.

Scalability and Migration

Each of the HP standard SAN topologies can be scaled incrementally to increase connectivity and overall capacity. You should always plan for expected future growth in your initial SAN design to minimize disruption when expanding capabilities and capacity over time. If you do exceed the capacity of a given topology, or find that data access needs have changed, it is possible to migrate one topology to another. Refer to Chapter 11, "Best Practices" for information about migrating topologies.

Table 16 lists the migration paths and the options for scalability for all topologies.

Table 16: Topology Migration & Scaling

SAN Topology	Migration	Scalability (For All Topologies)
Cascaded	Convert to Meshed, Ring or Tree	<ul style="list-style-type: none"> Increase the number of switches
Meshed	Convert to Ring, or Tree	<ul style="list-style-type: none"> Use higher port count switches
Ring	Convert to Meshed or Tree	<ul style="list-style-type: none"> Transition to a different topology
Tree	Add additional backbone switches	<ul style="list-style-type: none"> Deploy multiple fabrics

Custom-Designed SAN Topologies

The HP standard SAN topologies, or subsets of these topologies, as discussed in this chapter, can meet most SAN implementation requirements. There may be specific cases where HP standard topologies (or variants) do not meet your specific needs or requirements. In these cases, a custom SAN design can be created if the SAN design rules described in this document are strictly followed. Refer to Chapter 3, "SAN Fabric Design Rules", Chapter 4, "Heterogeneous SAN Platform and Storage System Rules", and Chapter 11, "Best Practices".

Fibre Channel Routing

The remainder of this chapter discusses Fibre Channel routing technology in a SAN environment.

Fibre Channel routing provides three basic features:

- A definition of the meaning and scope of SAN subnetworks
- A method for enabling communication between subnetworks while maintaining subnetwork independence
- A way to centralize management of subnetworks.

These features enable solutions to a wide range of issues that arise in constructing and managing SAN configurations. These issues include

- Interconnection of existing SANs
- Sharing devices between multiple SANs
- Geographically expanded configurations
- Configurations with many devices
- Heterogeneous installations
- Isolation of fault domains
- Centralized management
- Tape backup consolidation

The basic function of a Fibre Channel router is to connect multiple Fibre Channel switches together so that they can communicate. This is shown in Figure 21. Each subnetwork is a connected set of switches.

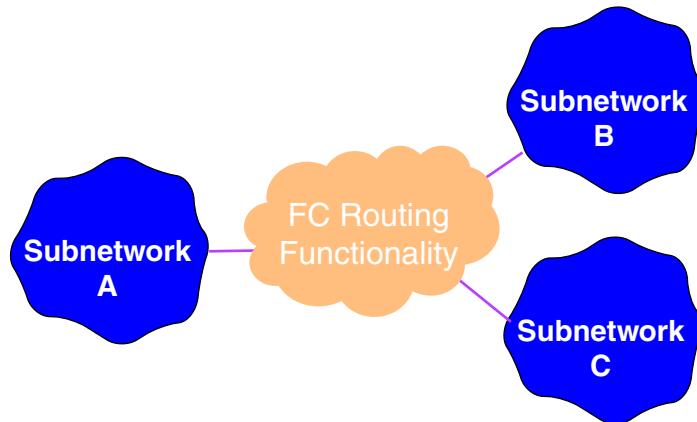


Figure 21: Basic Fibre Channel Routing Configuration

While Fibre Channel routing is generally similar to IP routing, there are significant differences in the way the routing technology works in the two cases.

The key driver for the development of Fibre Channel routing was the realization that the current method of scaling SANs is not practical in the long term. Successful scaling of traditional SANs requires scaling of the *distributed fabric services*. Fabric services provide functions that need to be coordinated between all the switches in a fabric. For example, the Simple Name Service provides a mapping between device names and their addresses in the network. To deliver a packet to the desired destination, every switch in the path from the

source to the destination must have this addressing information, and the information must be kept up to date over time and as the network configuration changes. This coordination is managed by distributing the service throughout all members of the network.

There are about a dozen fabric services altogether, including the Fabric Login Server, State Change Notification Server, Name/Directory Server, Zone Server, Key Server, Time Server, Alias Server, etc. Routing information maintained in the switches in a fabric is also considered a distributed fabric service.

Using traditional Fibre Channel terminology, each fabric is a network of connected switches that maintains a single unique set of fabric services. When two fabrics are connected in a traditional SAN installation, the fabrics merge their distributed services to obtain a new single unique set.

As fabrics become larger, the burden of supporting the fabric services becomes heavier. Additional resources are needed as the fabric grows, making it difficult to mix very small, inexpensive switches with large, resourceful switches. Without routing, fabric scaling is limited by the ability of the smallest switch in the system to participate in the distributed fabric services system.

The traditional definition of “fabric” must be refined when routing technology is used. This is discussed in more detail below.

Merging versus Dividing Fabrics

It's clear that the configuration shown in Figure 21 could be developed either by connecting several fabric subnetworks together or by dividing a single fabric into several subnetworks. Both approaches are used in Fibre Channel routing solutions.

In the case of B-Series products, multiple subnetworks are connected together. Subnetwork A is a connected set of one or more switches, as are Subnetwork B and Subnetwork C. Any B-Series switch product can be used in this type of configuration, and each subnetwork follows the traditional configuration rules for a fabric. A Multi-protocol Router is used to connect the fabrics together.

In the case of C-Series products, a given single fabric is divided into several subnetworks.

Subnetwork A is a set of switch ports on one or more switches, as are Subnetwork B and Subnetwork C. The routing function is distributed across all the switches in the original fabric, and there is no separate router hardware.

The two methods of using routing solve a similar set of problems. For example, either method can be used to connect a set of existing fabrics. Using a B-Series router, the existing fabrics are physically connected to the router and then the router is configured to provide the required communication. Using C-Series switches, the existing fabrics are physically connected together and then the routing function in the switches is configured as required.

Device sharing is defined in the routing configuration description. A device in Subnetwork A can be exposed to devices in Subnetwork B or Subnetwork C as required. For example, a centralized tape library could be configured so that servers in all three fabrics can communicate with it.

Product-Specific Differences

The B-Series and C-Series Fibre Channel switch product families offer slightly different routing solutions. The essential differences are described below, and additional product-specific information is available in the technical documentation for each product. Refer to the SAN Infrastructure web page at:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>

In the case of B-Series products, a distinct, separate hardware device, the Fibre Channel Multi-Protocol Router, is used to connect subnetworks together. Each subnetwork is made a collection of physical switches that forms a traditional Fibre Channel fabric, so the usual rules about fabrics apply. Within each fabric all the switches must run the same version of switch firmware for each switch model family and must have the same values for any variable settings (for example, R_A_TOV). Each fabric has a unique set of fabric services.

The router provides isolation between the fabrics, so two fabrics, for example Subnetwork A and Subnetwork B, could be configured with different firmware versions or with different variable settings.¹ This provides flexibility when connecting existing fabrics that might have different operating conditions or have different settings.

In B-Series configurations, switch ports can be grouped together into “LSANs.” An LSAN is similar to a traditional Fibre Channel zone, except that an LSAN can extend through a router to include devices in another fabric. The combined configuration, including the physical fabrics (that is, the subnetworks), the LSANs and the router, is called a “Meta-SAN.”

In the case of C-Series products, the routing function is contained with the switches and is enabled by a license. Since the complete collection of switches is one connected set of switches, all the switches must run the same version of switch firmware.

In C-Series configurations, switch ports can be grouped together into “VSANs.” A VSAN may extend across multiple switches. Each VSAN has a unique set of fabric services, so it is similar to a traditional Fibre Channel fabric in terms of management and independence. The variable fabric settings apply on a per-VSAN basis, so they may be different for each VSAN.

When discussing C-Series products, the subnetworks are called “fabrics,” the individual logical fabrics are called “VSANs” while the combined configuration is simply called a “SAN.”

Figure 22 and Figure 23 show the essential difference between configurations using B-Series and C-series products.

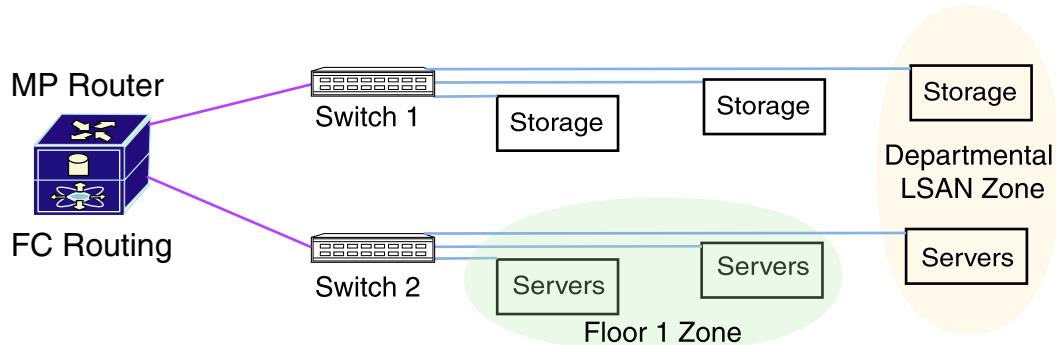


Figure 22: B-Series Routing

In the B-Series case, the LSANs are connected together at the router. An LSAN may consist of devices connected to the same switch (e.g. “Floor 2 LSAN”) or to multiple switches (e.g. “Departmental LSAN”).

1. Certain restrictions apply, as discussed in Chapter 3.

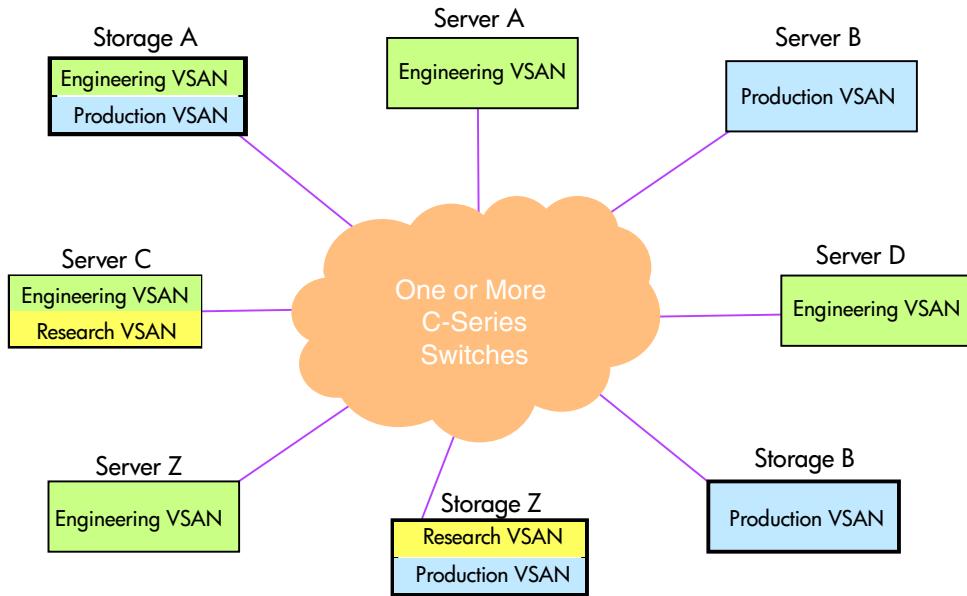


Figure 23: C-Series Routing

In the C-Series case, each VSAN may include devices that reside on a single switch or multiple switches in the SAN. Devices in different VSANs can communicate using the inter-VSAN routing feature. If using multiple switches, the switches are connected together in any conventional supported fabric configuration.

Routing Technology Overview

When you connect fabrics with Fibre Channel router technology, the routing function passes selected Fibre Channel frames² between the fabrics. The frame selection system allows the routing function to isolate the fabrics from each other from the viewpoint of the fabric services protocols. This allows the connection of multiple heterogeneous fabrics, provides the desired scaling, and enables fault isolation. The central location of the routing function makes it possible to provide a common management domain for all of the connected fabrics.

Network scaling is improved because if a fabric supports, for example, 1200³ ports, and we connect three of them together with a router, then we get a “3600 port SAN.” This sort of configuration *does not* provide a 3600 port network where any port can talk to any other port. What it does do is allow you to configure some subset of the ports, perhaps 1000, so they can communicate. A mapping table in the router describes the relationship between the ports that can communicate with each other.

This approach turns out to be satisfactory for many environments.

The advantage of this approach is that the system designer gains a considerably more flexible set of configuration rules, due to the “multiplier” effect of routing. For example, suppose that a stand-alone 1200 port fabric is supported. Then suppose that three fabrics are supported by a router with 1000 mapping table entries. The support rules for the combined configuration will be for $1200 \times 3 = 3600$ ports in total, with 1000 mappable ports.

The support tables in Chapter 3 reflect this viewpoint.

2. packets

3. Refer to the next chapter for the correct supported configuration sizes.

Traditional SAN Scaling

Using the conventional approach, ports are added to a SAN by adding switches to an existing SAN or by using higher port-count switches. Two limits to SAN size must be taken into consideration.

- Architectural limits on switch count and ports on a switch.
- Support limits.

The Fibre Channel architecture allows for a maximum of 239 switches in one fabric, and a maximum of 256 ports per switch.⁴ The practical limit to SAN size is quite a bit smaller than these theoretical limits. These practical configuration limits are determined by lab testing and by understanding the internal design of the SAN infrastructure components.

Configuration rules define how many switches may be used and how many ports are supported. These rules offer limits on both dimensions so that if you select large switches, you are likely to run into the port count limit first, while if you use small switches you may run into the switch count limit first. HP has made large strides in support for larger and larger SANs, but some particularly large customers are asking for solutions that put pressure on what is known to work properly.

Today's standard support limits are generous enough to satisfy the requirements of practically all users.

Two limits are listed for each Fibre Channel SAN switch product line,⁵ and a SAN designer must verify that a new design or a modification to an existing SAN complies with these limits.⁶

- Total port count. A maximum number of ports may be present in a fabric. This limit includes the switch ports used to connect hosts and storage to the storage network and also includes the switch ports used for ISLs that connect the switches together. The user port count in a given network will be somewhat lower than the total port count, depending on the network topology chosen.
- Total switch count. A maximum number of switches may be present in a fabric. This limit insures that the distributed fabric services algorithms in the network are not exceeded.

By requiring that a SAN design must simultaneously meet both the total port count limit and the total switch count limit, several undesirable configurations are avoided. First, the SAN designer understands the basic limit of network size, in both dimensions. Second, the designer understands that using many small switches to reach a high total port count number is not acceptable if it exceeds the total switch count limit. Third, using a moderate number of extremely large switches is not acceptable if it causes the design to exceed the total port count limit. By quoting the limits in this fashion, HP makes the design of a large configuration as straightforward as possible.

“SAN” versus “Fabric”

Early in this chapter definitions are provided for the terms “SAN” and “fabric.” When discussing networks that include routers, these terms must be viewed from a slightly different perspective. As mentioned previously, a fabric is a set of switches connected together into a network, all supported by a single set of fabric services. The fabric services are supporting functions that are used to manage device names and addresses, timestamps, and other utility functionality used by the switches themselves.

4. Refer to the Fibre Channel standards for more information.

5. B-Series, C-Series, M-Series

6. Other limits are also applicable, including hop count, link distances, etc.

A SAN is a connection of servers and storage devices, plus their interconnecting fabrics. Most large SANs have two fabrics in order to maintain communication in the case of the failure of one of the fabrics.

Using these definitions, a set of switches may be physically connected together into a single fabric, or a set of switches may be physically connected together but partitioned into LSANs or VSANs that are each fabrics in themselves. This clarification of the definition of “fabric” allows a single interconnected network to contain multiple fabrics, and it allows the construction of networks where devices on different fabrics can communicate while maintaining a large degree of independence.

Larger SANs

Given practical storage network configuration limits, how does HP support ever-larger SANs? To begin with, by providing larger basic storage network configurations. The limits for total port count and total switch count are increased on a frequent basis to reflect customer demand for large configurations.

Another dimension to scaling a storage network is by adding fabrics. A minimal storage network is a single switch with servers and disk arrays attached to it. This is a single fabric, and adding more switches by connecting them to the first switch makes that single fabric larger. A dual fabric SAN can be built by using two unconnected switches with servers with dual HBAs and storage devices with dual ports. The principal reason for choosing a dual fabric SAN is that if one fabric fails then the other fabric provides redundancy. A side effect of choosing a dual fabric SAN is that there are more ports in the complete solution.

A question may arise at this point about how large storage networks should eventually become, and how storage networks are similar to and different from general purpose corporate networks. HP views storage networks as *complementary* to general purpose networks. General purpose networks serve primarily as a method of obtaining connectivity so that many servers can communicate with many other servers. This is the basis of the Internet and of the majority of office LANs, MANs, and WANs.

Storage networks are different from general purpose networks because they require performance and availability that is far in excess of usual corporate networks. A general purpose LAN typically has a single connection to each server. A storage network generally has a pair of redundant connections, and this redundancy is carried throughout the entire storage network. A general purpose LAN typically has comparatively low performance: Most office environments today use 100 Mbps infrastructure. Storage networks require high bandwidth: Most SANs today use 2 Gbps Fibre Channel infrastructure. Storage networks require high performance even as they extend to great distances through large networks, while typical LANs serving thousands of desktop users may use 45 Mbps DS3 connections between sites.

The vision of future storage networks may be summarized in two points.

- Future SANs will continue to expand in all dimensions
- Customers will choose to maintain separate and parallel deployments of storage networks and general purpose networks for the foreseeable future.

How Does Fibre Channel Routing Work?

The basic idea of routing is to join two or more independent network fabrics together in a way that enables communication between ports on the fabrics but maintains a degree of isolation between them.

Using this neutral terminology,⁷ a SAN may be defined as a storage system consisting of storage elements, storage devices, computer systems, and/or appliances, plus all control software, communicating over a network. A SAN may have multiple fabrics where each fabric consists of one or more switches, as discussed above.

A router is a network component that forwards frames from one fabric to another, thus connecting the two fabrics. Routers interact with switches in a network and cooperate to forward frames to the correct destination. A router reads the fabric address information in each frame that it receives, then uses a routing table to determine the destination fabric and address. The router then transmits the frame into the destination fabric.

The result of joining fabrics together using a router is a single extended fabric. Devices in each fabric may be exported to another fabric through the router. When a router maps an address in a given fabric to an address in another fabric, the mapped address in the latter fabric is the imported address. Similarly, when a router maps an address in a given fabric to an address in another fabric, the original address is the exported address.

An imported device is a device as seen in a fabric when using its imported address. An exported device is a device as seen in a fabric when using its address.

Fibre Channel Addresses

Every device in a network has a name and an address. The name is a human-readable descriptor, while the address is embedded in the packet headers and is used by the network switches to direct packets to the device.

Fibre Channel defines the device name and address so that the name is unique and fixed, while the address varies. This is opposite to the way addressing is done in the IP over Ethernet world, and this difference causes some difficulty in describing routing.

In Fibre Channel:

- A recognized naming authority assigns the *device name* so that it is a worldwide unique identifier. This is called the “*worldwide name*” or “WWN” of the device. The network manager uses the WWN to unambiguously refer to each device during the setup phase of the SAN. WWNs are used to define which devices are to be exported from one fabric to another and to assign devices to zones.
- The network switches assign the *device address* during the fabric initialization process, so that each device address corresponds to the switch port to which it is attached. This is called the “*N_Port ID*” of the device. The network uses the Destination N_Port ID found in the header of each frame to direct the frame to the proper destination edge port of the fabric. Since the fabric assigns the N_Port ID, two fabrics may use the same N_Port ID to refer to different devices.

A Fibre Channel switch uses a name server to maintain the relationship between WWNs and N_Port IDs. Numerous control messages are used in Fibre Channel fabrics to control the behavior of the name server and other related fabric services needed for network operation.

7. This terminology is based on the SNIA SAN dictionary, copyright used with permission, and the enhancements needed to describe routing technology have been proposed for addition to the dictionary.

Address Handling when Routing

If a Fibre Channel router is used to connect two fabrics, one important task that it must accomplish is to map the N_Port IDs between the two fabrics. This is done by using a routing table. Unlike IP,⁸ the Fibre Channel protocol does not provide an automated way to populate the routing table. The network manager defines the relationships between WWNs and N_Port IDs to reflect the way that ports are to be presented in fabrics.

Since each fabric has a separate set of fabric services, the router provides isolation between the fabric services instantiations in each fabric. This has the desirable side effect of reducing the amount of traffic that passes through the router, and in the case of the fabric-wide failure of a fabric service provider, the outage is limited to that fabric only.

Using this method the size of the network is limited by the size of the N_Port ID address space, which is quite large, and by the size of the routing table.

In Fibre Channel, each frame of data sent between an originating and a destination port has a CRC. The CRC covers both the data and the frame header. In a Fibre Channel switch, as soon as the destination address field in the header has been examined the frame can be sent out on the appropriate switch output port. This “cut-through” approach minimizes latency in the network, and is a significant feature of Fibre Channel compared to many other network protocols.

When a frame passes through a router, the address field in the header is changed so that the destination fabric can deliver the frame to the right destination port. This means that the CRC must be recalculated. To do this, the router holds the entire frame in a buffer.

8. IP/Ethernet routers use ARP, the Address Resolution Protocol, to populate their tables.

Supported Configurations

When routing capability is to be included in a SAN design, additional configuration rules must be followed. This section describes the general design considerations, and the next chapter defines the detailed configuration rules and limits.

The typical configuration with routing has a router connected to two or more completely separate fabrics, as shown in Figure 24. Using B-Series components a router is required, while using C-Series components the fabrics are connected directly to each other.

Routers Connecting Core-Edge Fabrics

If the fabrics have core-edge topologies, it will be found advantageous to connect cores together, as shown below.

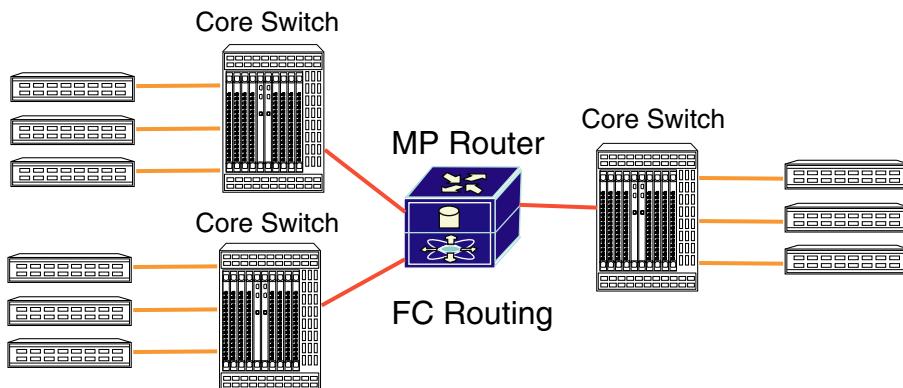


Figure 24: Fabrics Connected at Cores using the B-series MP Router

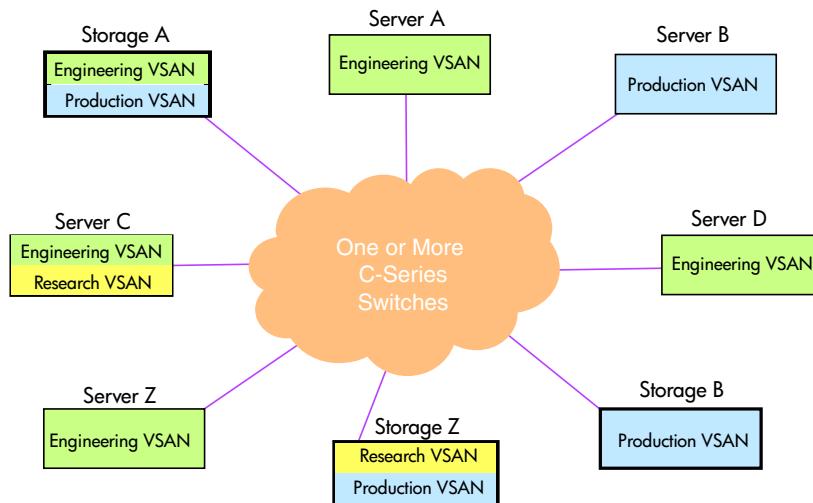


Figure 25: VSANs Connecting Core Switches

Routers Connecting Fabrics at a Distance Through an IP Network

When used to connect fabrics located at a distance from each other through IP, routers may be incorporated into the fabric as required. However, since the routers communicate using the FCIP protocol instead of normal Fibre Channel, they must be installed in pairs as shown in Figure 26.

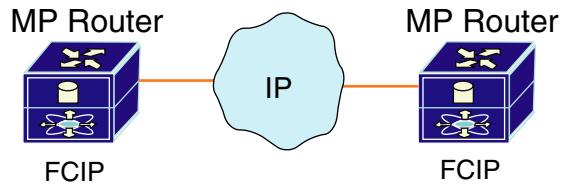


Figure 26: Routers Connecting Fabrics at a Distance

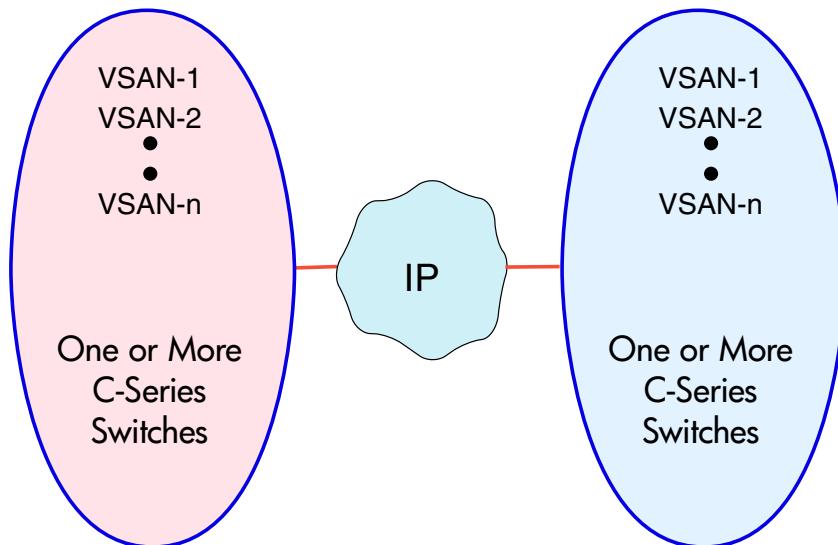


Figure 27: VSANs Connecting over FCIP

Router High Availability Configurations

When used in high availability configurations, routers may be used in pairs to maintain isolation between fabrics. All of the configurations shown below are valid. The top configuration has servers and storage connected using a pair of completely redundant fabrics in a level 4, NSPOF configuration (refer to [Data Availability in a SAN](#), page 63 for information on high availability levels). The second configuration shows routers connected together to provide failover between fabrics in case of the failure of one router. The third configuration shows routers cross-wired to provide full fabric connectivity in case of a failure of one router. The bottom configuration shows independent fabrics for the server and storage, connected using a single router. Availability in this configuration can be increased by adding additional ISL connections between the MP Router and the fabrics.

For simplicity, the configurations below show a small number of fabrics connected to each router. The concepts presented hold true for larger configurations with a higher number of fabrics connected per router, and a higher number of routers per Meta SAN. Refer to Chapter 3, "[SAN Fabric Design Rules](#)" for scalability rules including the maximum number of fabrics and MP Routers per routed fabric when using the B-Series MP Router.

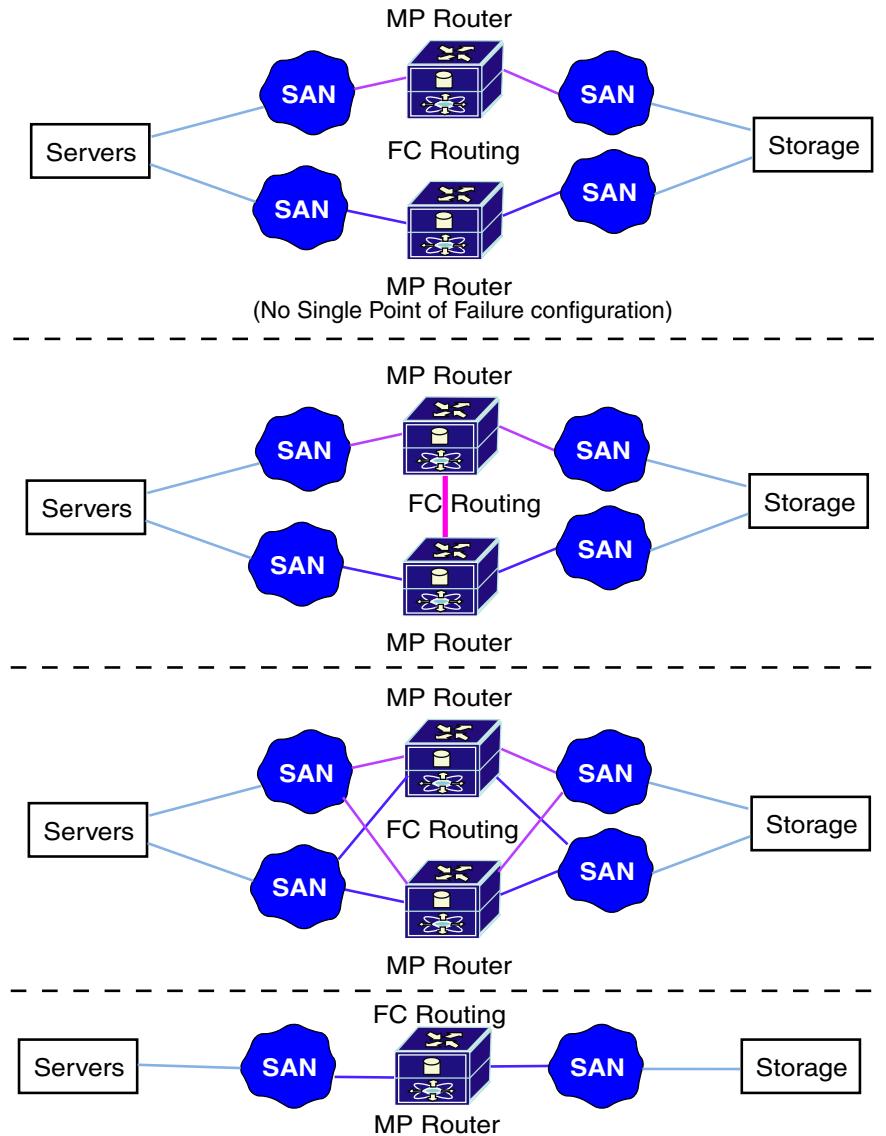


Figure 28: High Availability Configurations with Routers

Unsupported Configurations

If fabrics are connected by a router, there cannot be a regular ISL link between the fabrics in parallel with the router. Such a link provides a direct path between the fabrics, which makes them into a single fabric and bypasses the isolation and mapping capability of the router..

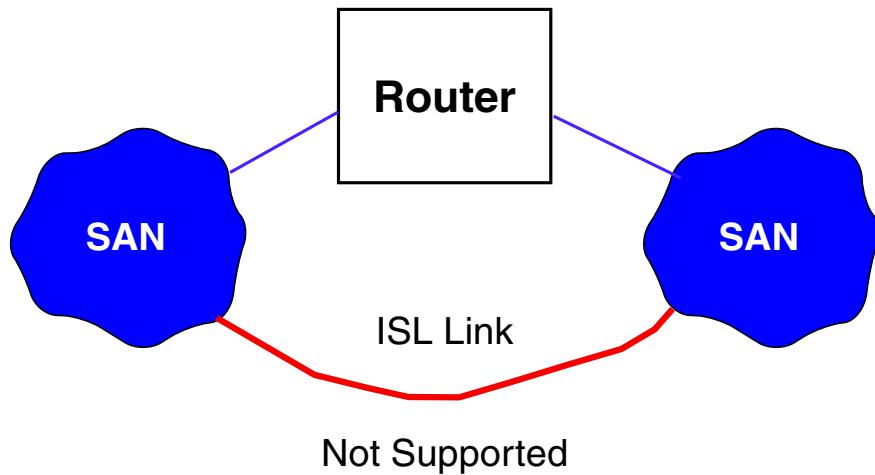


Figure 29: Unsupported Configuration

Refer to Chapter 3, "[SAN Fabric Design Rules](#)" for the HP supported MP Router configuration rules. All MP Router configurations must conform to the configuration rules listed in Chapter 3.

Summary

Fibre Channel routing brings the ability to construct and manage *much larger storage network configurations*. This is accomplished by using three basic capabilities of routed fabrics:

- Scaling to larger networks.
- Isolation of subnetworks to confine problems.
- Common management of subnetworks.

The routing technology, whether implemented in a separate Multi-Protocol Router or in the Fibre Channel switches, provides the essential ability to scale to larger networks.

Isolation is provided by the routing technology because each fabric has a separate, independent copy of the fabric services.

The subnetworks, or fabrics, are managed using a single point of control, which reduces the number of management interfaces from one per fabric to one per SAN.⁹

9. Or two per SAN if a pair of redundant fabrics is used.

MP Router Use Case Configurations

SAN Island Consolidation/SAN Scaling

One of the primary uses of the MP Router Fibre Channel Routing feature is to consolidate multiple SAN islands into a Meta SAN, in effect, creating a modular SAN design. This simplifies scalability by allowing you to scale your SAN for the purpose of adding and sharing additional devices selectively without the need to fully merge all fabrics. Typically, access to devices in different fabrics is only required for a subset of all devices in each fabric, and on a limited time basis such as for data migration and storage consolidation. Using the MP Router also eliminates the need to physically move equipment from one location to another to provide access. Connecting multiple fabrics to the MP Router instantly allows for the sharing of devices located anywhere in the Meta SAN.

Since the MP Router does not actually merge fabrics, existing zoning definitions and assigned domains can be used without modification, even in cases where the same zoning definitions and domain IDs have been used for multiple fabrics connected to the same MP Router. Individual fabrics within a Meta SAN can be scaled without affecting or limiting scalability in other fabrics.

From a management perspective, management of multiple SANs can be centralized and consolidated to one Meta SAN or partitioned into different administrative domains as required. Whatever approach is taken, HP strongly recommends the use of Fabric Manager to simplify management procedures when implementing an MP Router based Meta SAN.

Figure 30 shows a typical configuration using the MP Router with Fiber Channel Routing for SAN island consolidation.

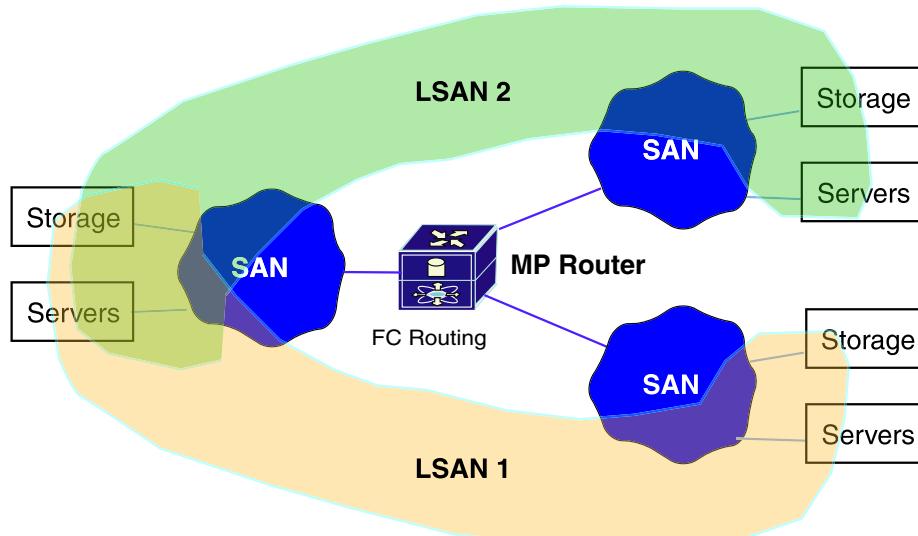


Figure 30: Consolidating SAN islands with the MP Router

Integration of Fibre Channel Routing and FCIP

The combination of Fibre Channel routing and FCIP available in the HP StorageWorks MP Router provides an attractive solution for disaster tolerant applications such as Continuous Access for HP storage arrays. In typical Continuous Access configurations without the MP Router, local and remote fabrics are completely merged when connected through FCIP since this effectively creates an ISL connection through an IP network. By using the MP Router

Fibre Channel Routing and FCIP functions together, you can connect the local and remote fabrics together without fully merging them. LSANs can be created to allow selective connectivity of the local and remote storage arrays and associated servers without the need to merge both fabrics.

Figure 31 shows a typical No Single Point of Failure (NSPF) configuration using the MP Router for Fiber Channel Routing and FCIP with Continuous Access.

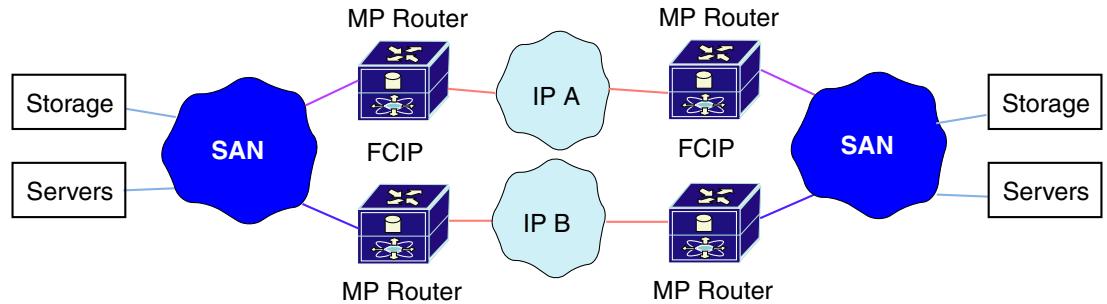


Figure 31: Continuous Access NSPF configuration using the MP Router for FCIP/FC routing

Tape Backup Consolidation

Just as the initial implementation of Fibre Channel attached tape devices and libraries encouraged consolidation of backup functions in a SAN, the MP Router provides a means to extend tape consolidation across multiple fabrics. This higher level of consolidation enables tape backup for devices residing in fabrics without tapes. Tape libraries and backup operations can be centralized and shared across multiple fabrics in a Meta SAN without the need to merge individual fabrics, reducing overall equipment and management costs.

Figure 32 shows a typical configuration for using the MP Router to consolidate tape backup for several SANs onto a single SAN. Refer to the HP StorageWorks Enterprise Backup Solutions Design Guide and EBS compatibility matrix for information on supported HP tape products and backup applications.

<http://h10025.www1.hp.com/ewfm/wc/manualCategory?lc=en&cc=us&product=406722>

<http://h18004.www1.hp.com/products/storageworks/ebslegacymatrices.html>

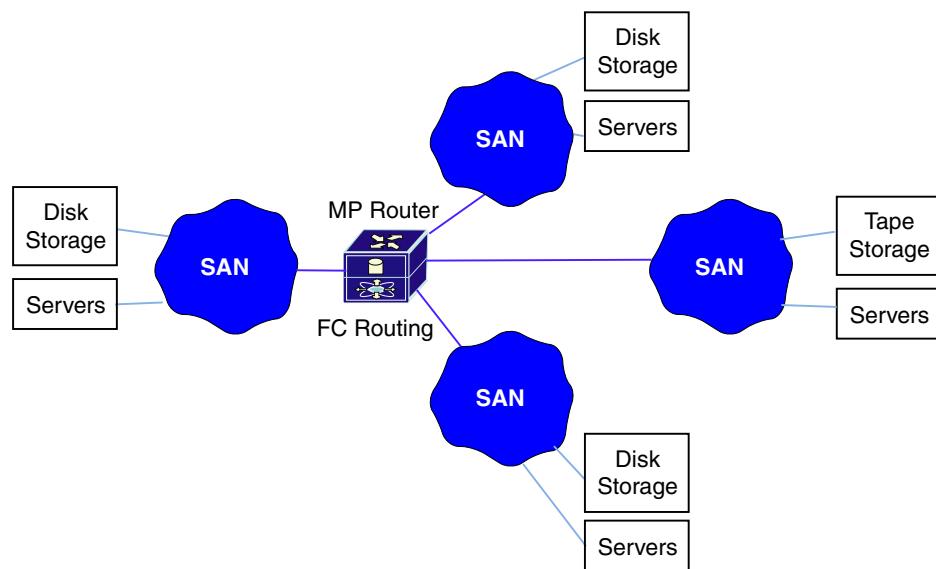


Figure 32: Tape backup consolidation

